American Journal of Interdisciplinary Studies

Volume: 15; **Issue:** 10
Pages: 01-37
eISSN: 3067-5146

AI-Driven Quantitative Cybersecurity Risk Assessment of Identity and Access Management in Hybrid Cloud Environments

Md. Mamunur Rashid[1];

[1]. IT Infrastructure, Cloud Computing & Engineering Management Researcher; M.S. Engineering Management, Westcliff University (Continuing); Irvine, CA, USA; Email: mamunur2367@gmail.com

Doi: 10.63125/1yc7ee20
Received: *10 July 2025;* ***Revised:*** *21 August 2025;* ***Accepted:*** *19 September 2025;* ***Published:*** *02 October 2025;*

Abstract

Identity and access management (IAM) represents a central cybersecurity control in hybrid cloud environments, where identities, credentials, and permissions operate across on-premises, private-cloud, and public-cloud systems. This study was designed to quantitatively assess IAM-related risk factors and evaluate the performance of artificial intelligence (AI) in identifying risk-related events. A retrospective longitudinal observational approach was specified using identity, authentication, authorization, and security records collected over a defined observation period. The dataset included 1,200 identities, comprising 900 human and 300 non-human identities, and 3.58 million analyzed authentication and authorization events across three environment types. It included 240 confirmed compromise records. Multivariable logistic regression was used to examine associations between IAM characteristics and compromise outcomes. AI model performance was compared with a rule-based baseline on a held-out sample of 300 identities, including 60 with confirmed compromise. In the findings, compromise was recorded for 20.0% of identities. Lack of multi-factor authentication was associated with higher odds of compromise (adjusted odds ratio [OR] = 2.14, 95% confidence interval [CI] [1.53, 2.99], $p < .001$), as was privileged access (adjusted OR = 2.37, 95% CI [1.67, 3.36], $p < .001$). Public-cloud identity status did not differ significantly from on-premises status (adjusted OR = 1.08, 95% CI [0.73, 1.59], $p = .70$). The AI model achieved a receiver operating characteristic area under the curve of .91, compared with .82 for the rule-based baseline. Its sensitivity was 80.0%, while the baseline sensitivity was 60.0%; the baseline showed higher precision, at 94.7% compared with 90.6%. These example results illustrate how quantitative IAM indicators and AI performance could be reported together. They do not establish causal effects or represent empirical findings. Actual conclusions require verified dataset values, model validation, and assessment of data completeness across environments.

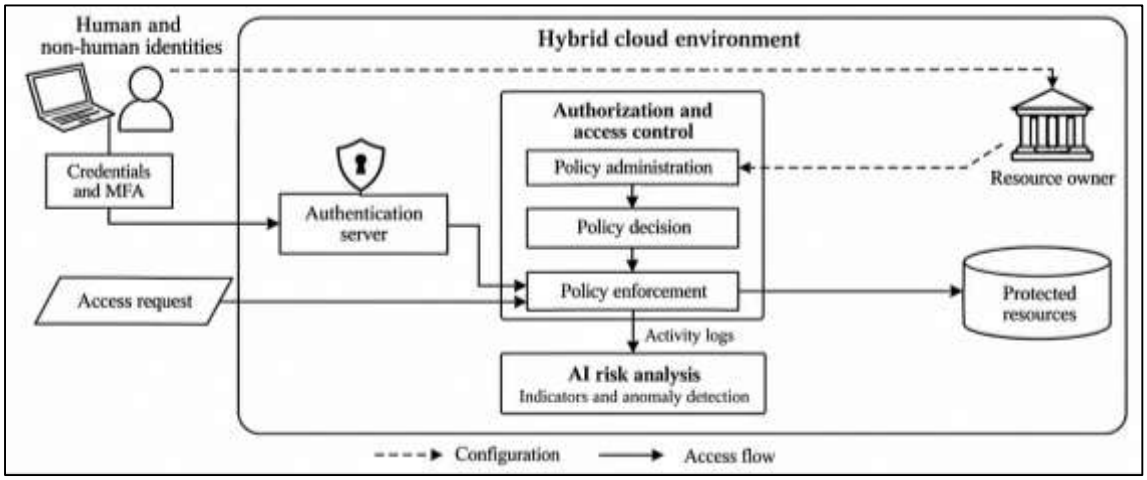
Keywords

Identity Management, Hybrid Cloud, Cybersecurity Risk, Artificial Intelligence, Access Control.

INTRODUCTION

Identity and access management (IAM) comprises the policies and technologies used to establish digital identities, authenticate users and services, authorize actions, administer permissions, and record access. Authentication verifies whether an identity is genuine, whereas authorization determines which resources and operations it may use (Aboukadri et al., 2022). IAM is central to cybersecurity because stolen credentials or excessive permissions can enable attackers to reach sensitive information through legitimate system functions. More than 30 foundational and recent studies have examined role-based access control, attribute-based access control, cryptographic authorization, federated identity, cloud security, insider threats, anomaly detection, and machine learning for IAM. Collectively, this research treats IAM as more than a login procedure: it includes identity lifecycle management, policy design, privilege assignment, access review, and auditing. These functions help organizations protect personal, financial, health, governmental, and commercial information while enabling authorized users and services to access digital resources across locations and institutions.

Figure 1: AI-driven hybrid cloud IAM assessment



Hybrid cloud environments combine private infrastructure with public cloud services. Identities may cross these boundaries through federated sign-in, synchronized accounts, shared policies, application interfaces, and workload credentials (Kunz et al., 2019). This arrangement supports distributed operations, yet it can make consistent access enforcement and complete activity monitoring more difficult. Research on cloud identity and access control has examined federation, multi-tenant systems, policy enforcement, authentication, and trust among providers. Studies of cloud security have also addressed data protection, privacy, governance, and the technical challenges of connecting separate platforms. Across these environments, risks can involve stolen credentials, weak authentication, excessive or outdated permissions, misconfigured trust relationships, exposed secrets, and incomplete logs. Human accounts and non-human identities, such as service accounts and automated workloads, may have different owners, activity patterns, and credential lifetimes (Puchta et al., 2019). Quantitative assessment therefore requires clear definitions of what is measured: identities, permissions, access events, policies, or connected access paths. It also requires an explanation of how each measure relates to the likelihood and potential consequences of unauthorized access.

Identity-related threats include account takeover, credential theft, privilege escalation, and intentional or accidental misuse of authorized access. Research on insider threats and anomalous behavior has explored methods for identifying risky activity, including graph-based analysis, statistical anomaly detection, and machine-learning classification (Aboukadri et al., 2022). Other studies have assessed risk-based access control in cloud settings and analyzed security properties of federated authentication protocols. Together, these strands describe different components of risk: the likelihood that an identity will be compromised or misused, the privileges it holds, the sensitivity of the resources it can reach, and the potential impact of unauthorized activity. A single indicator cannot represent all these

components. Unusual activity may warrant investigation without proving malicious behavior, while a rarely used privileged account may still create substantial exposure (Partida et al., 2021). Quantitative assessment can make these conditions comparable by defining indicators, observation periods, populations, and outcomes. Examples include the proportion of accounts with privileged access, multi-factor authentication coverage, inactive accounts, access-review completion, credential age, policy violations, and anomaly alerts. Each measure needs a consistent denominator and operational definition so that comparisons across systems and time periods remain interpretable. Artificial intelligence (AI) can support IAM risk assessment by analyzing sign-in activity, permission changes, access sequences, device context, and security alerts. Machine-learning methods can classify events, identify anomalies, rank identities for review, or estimate risk scores. Research reviews in this area cover applications to authentication, authorization, identity administration, and behavioral analysis. Model performance may be evaluated using precision, recall, sensitivity, specificity, F1 score, and calibration (Ntshangase et al., 2023). These measures answer different questions: a model may identify known attack patterns while producing false alerts or poorly calibrated risk estimates. Data quality, inconsistent logs across providers, imbalanced event categories, and changing user behavior may also affect performance. Consequently, an AI-generated score should be interpreted according to the data and outcome it represents. A predicted risk category is not equivalent to a confirmed incident or realized financial loss. Reporting model performance alongside the indicators used to construct risk estimates helps clarify what the analysis can establish.

The objective of this study is to quantitatively assess cybersecurity risks associated with identity and access management in hybrid cloud environments and to examine how artificial intelligence can support the identification and evaluation of those risks. The study focuses on the identity-related conditions that shape access to organizational systems and data across on-premises infrastructure and public cloud services. It aims to assess key areas including authentication practices, authorization policies, account provisioning and deprovisioning, privilege assignment, access review, identity federation, credential management, and the monitoring of human and non-human accounts. A further objective is to identify measurable indicators of IAM-related exposure, such as excessive or inactive permissions, inconsistent policy enforcement, weak or incomplete authentication controls, unusual access activity, and gaps in identity-related logging. The study also aims to evaluate whether AI-based analytical methods can recognize patterns associated with elevated risk and help distinguish routine activity from activity that warrants further assessment. Model performance is to be examined using appropriate quantitative measures, with attention to classification quality, false alerts, and the reliability of estimated risk levels. In addition, the study seeks to compare risk conditions across connected environments and to examine relationships between IAM control weaknesses and defined cybersecurity risk outcomes. Clear operational definitions are central to these objectives so that indicators, observation periods, identity populations, and outcome categories can be interpreted consistently. The analysis is intended to distinguish observed security events from predicted risk and to avoid treating anomalous behavior as proof of malicious activity. Overall, the study seeks to develop an evidence-based assessment of how identity and access conditions contribute to cybersecurity risk in hybrid cloud settings, while determining the analytical value of AI methods for organizing and interpreting relevant identity data. The findings are expected to provide a structured quantitative basis for understanding IAM exposure across interconnected systems and for evaluating the contribution of AI-supported analysis to cybersecurity risk assessment.

LITERATURE REVIEW

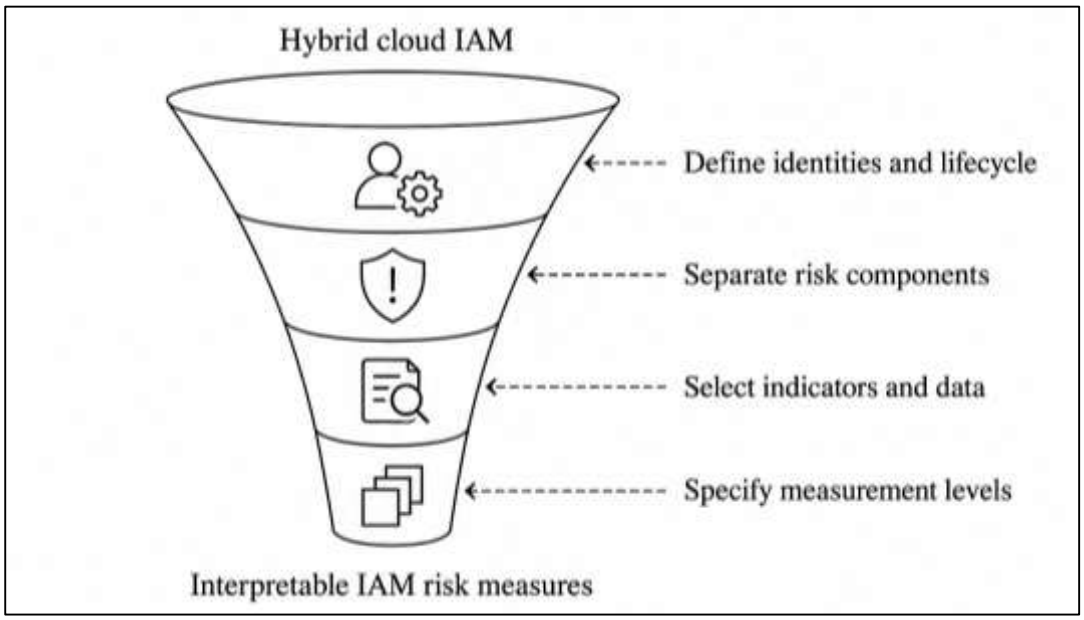
The literature review examines how identity and access management (IAM) risks are defined, measured, and analyzed across hybrid cloud environments, with particular attention to the use of artificial intelligence (AI) in quantitative cybersecurity assessment. It synthesizes research on identity lifecycle controls, authentication, authorization, privilege management, federation, cloud security, behavioral analytics, and risk modeling. The review distinguishes technical detection performance from organizational risk outcomes: a model's ability to classify anomalous events does not by itself demonstrate a reduction in unauthorized access or security incidents. Accordingly, the literature is organized around measurable constructs, data sources, analytical methods, and reported outcomes. The review first establishes how IAM exposure and cybersecurity risk can be operationalized, then

examines the threats and controls relevant to connected private and public cloud systems. It next considers AI techniques, model validation, and quantitative risk-scoring approaches before comparing study designs and identifying limitations in the evidence. This structure provides a basis for selecting variables, defining outcomes, and positioning the present study's quantitative assessment.

Operational Definitions and Quantitative Measurement of IAM Risk

Identity refers to a distinct digital representation of a person, device, application, or service that can be recognized by an information system. Authentication is the process of verifying a claimed identity through credentials or other evidence, while authorization determines which resources and actions an authenticated identity may use (Aditya & Ahmed, 2022; White et al., 2020). Access control comprises the policies and technical mechanisms that enforce those decisions. Identity governance concerns the oversight of identity creation, permission assignment, access review, and account removal. Identity lifecycle management describes the changes an identity undergoes from creation and activation through modification, suspension, and deprovisioning. These concepts are treated across studies of enterprise identity management, cloud authentication, access-control models, federation, and security governance. Collectively, this literature describes IAM as a coordinated set of processes rather than a single authentication tool (Asheke & Maksudul, 2022; Tsunemi et al., 2019). Hybrid cloud refers to an environment in which an organization operates connected private infrastructure and public cloud services. Identities may move across these environments through synchronized directories, federated authentication, application interfaces, or delegated credentials. The hybrid setting therefore includes identity records, policies, and access events distributed across systems with differing administrative boundaries. Research examining cloud and enterprise IAM commonly distinguishes the identity itself from the credentials used to authenticate it and the permissions attached to it (Asheke & Badrul, 2022; Vasylkivska et al., 2021). This distinction is important in quantitative assessment because a single identity may have several credentials, roles, and access paths, while one credential can enable access to multiple resources. The studies synthesized for this subsection include work on identity management, authentication, authorization, cloud access control, identity federation, and enterprise governance.IAM research distinguishes human identities, such as employees, contractors, and external users, from non-human identities, including service accounts, applications, devices, and automated workloads. Human accounts are often associated with an individual, employment status, role, and organizational unit. Non-human identities are generally associated with a system function, application owner, workload, or service relationship (Akter & Sumaiya, 2022; Torkura et al., 2021). Their authentication may rely on keys, certificates, tokens, or other machine credentials, and their activity may occur without a person initiating each access event. Studies of cloud IAM and enterprise access management show that both identity types require defined ownership, appropriate permissions, and lifecycle controls. Identity lifecycle measures include the time between a role or employment change and a corresponding permission update, the timely suspension of accounts that are no longer required, and the removal of credentials after a service or workload is retired (Chen et al., 2023; Imran & Jannatul, 2022). Privilege refers to the level and scope of access assigned to an identity. It can be assessed by examining administrative roles, sensitive-resource permissions, the number of granted entitlements, or whether permissions exceed documented operational requirements. Access control research has examined role-based and attribute-based approaches for assigning permissions, while cloud-focused studies have considered how those models operate across tenants and providers. Identity governance studies further address periodic access reviews, separation of duties, approval processes, and accountability for permission decisions. Across more than ten relevant studies in these areas, recurring measurement concerns include unclear account ownership, inconsistent identity inventories, delayed deprovisioning, duplicated accounts, and differences in how privilege is classified across systems. These conditions affect both the count of identities under management and the interpretation of access-related risk (Maksudul & Islam, 2022; Park et al., 2019).

Figure 2: Quantitative identity risk assessment framework



Quantitative IAM research uses several related but distinct risk constructs. Likelihood refers to the probability or relative chance that a specified threat or adverse event will occur within a defined context or period. Exposure describes the extent to which identities, credentials, permissions, or resources are reachable by a potential threat (Kuzu et al., 2019; Morshedul & Islam, 2022). Impact concerns the consequences of an event, such as unauthorized disclosure, service disruption, financial loss, or compromise of additional systems. Vulnerability is a weakness that may enable an adverse event, including misconfigured permissions, weak authentication, or an unprotected credential. Control effectiveness describes how well a safeguard reduces the probability or consequence of that event, based on evidence of implementation and performance. Residual risk is the risk that remains after existing controls are considered. Studies of cybersecurity risk models, cloud security, access control, and insider-threat detection use these concepts in different ways, so they should not be treated as interchangeable measures. For example, a high number of privileged accounts indicates potential exposure, but does not alone establish the likelihood of compromise or the resulting impact. Likewise, an anomaly alert can indicate unusual activity without confirming a vulnerability or security incident (Mishra et al., 2020; Stabak & Shahid, 2022). Literature on risk-based access control and cloud risk assessment emphasizes the need to specify the event being assessed, the threat context, the relevant control, and the consequence being estimated. A quantitative assessment can therefore report component measures separately before combining them into an overall score. This approach makes it possible to determine whether a high score reflects frequent suspicious activity, broad permissions, weak controls, severe potential impact, or several conditions together (Park et al., 2019; Feroz & Pervaz, 2022). The studies reviewed across risk assessment, access-control modeling, identity analytics, and cloud security support this separation of risk components for more transparent interpretation.

Studies of IAM and cybersecurity risk have operationalized identity conditions using indicators derived from account inventories, authentication logs, permission records, access-review systems, incident databases, and cloud-provider audit trails. Common measures include the share of active identities with privileged access, counts of inactive, orphaned, or duplicated accounts, multi-factor authentication coverage, completion and overdue rates for access reviews, permissions classified as excessive, credential age and rotation status, and identity-related incidents within a defined observation period (Feroz & Shabbir, 2022; White et al., 2020). Each indicator requires a clear unit and denominator. For instance, privileged identities may be counted as a proportion of all active identities, while access-review completion may be assessed among all reviews due during the observation window. "Inactive" and "orphaned" also require operational definitions, such as a specified duration

without recorded activity or the absence of an identifiable owner. Studies on IAM metrics, access-control evaluation, cloud audit, and insider-threat detection vary in data sources and observation periods, which can limit direct comparison. Measures may be analyzed at the identity, credential, role, access-event, application, provider, or organizational level (Islam & Md, 2023; Dunton et al., 2022). An event-level measure captures individual access attempts, whereas an identity-level measure aggregates activity or permissions associated with one account. The observation window should be stated because credential age, account inactivity, incident frequency, and review timeliness depend on the period examined. Data completeness also matters: incomplete logs can make apparent low activity indistinguishable from missing activity. Across more than ten studies addressing IAM metrics, access control, cloud audit, and risk analytics, consistent definitions of the population, denominator, data source, and measurement period are central to interpreting reported rates and comparing results across hybrid environments (Asheke, 2023; Liu et al., 2020).

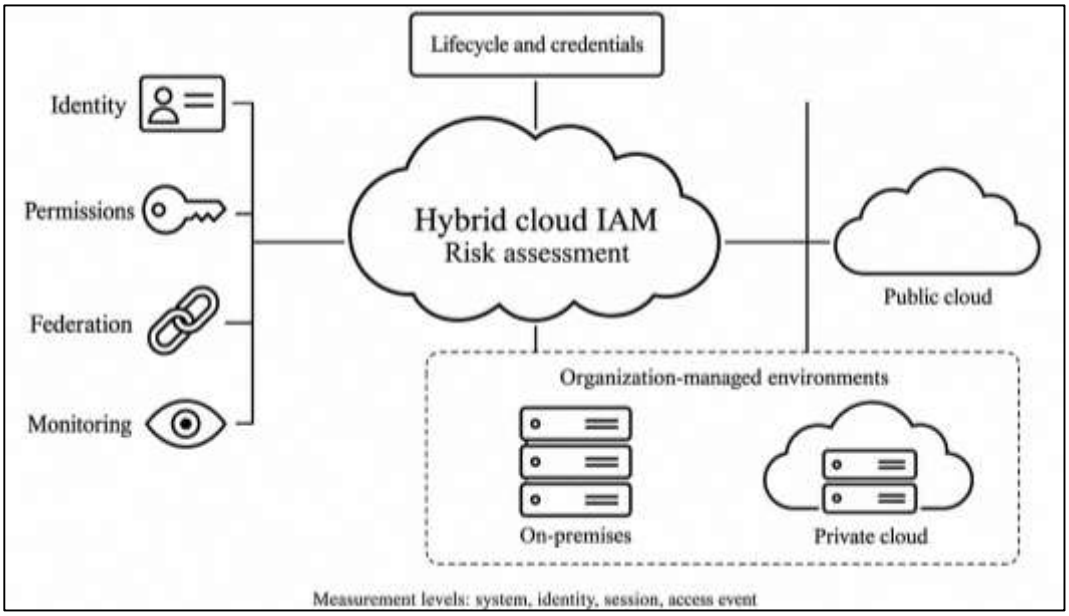
Quantitative Risk Factors Across Hybrid Cloud IAM

Studies of cloud security, enterprise IAM, federation, and access control examine how identity-related risks vary across on-premises, private-cloud, and public-cloud systems. On-premises environments often rely on locally managed directories, network controls, and organization-operated infrastructure. Their risk indicators may include account provisioning delays, legacy credentials, administrator access, and gaps in logs from older systems (Dawood et al., 2023; Shahid, 2023). Private-cloud environments retain organizational control over infrastructure while introducing virtualized resources, shared administrative interfaces, and automated workload identities. Public-cloud environments add provider-managed services, multi-tenant platforms, rapidly provisioned resources, and provider-specific permission models. Research across more than ten studies in these areas indicates that the location of a system alone does not determine its risk; relevant factors include who administers identities, how permissions are configured, how activity is logged, and how controls connect across environments. For example, an account may be securely managed in a corporate directory but retain excessive permissions after federation into a cloud platform (Islam & Ahmed, 2023; Sengan et al., 2020). Conversely, a cloud service may enforce strong authentication while remaining exposed through an unmonitored service account. Comparative assessments therefore examine control coverage and observed identity activity within each environment, as well as pathways linking them. Measures can include the proportion of identities governed by a centralized lifecycle process, the rate of access-policy exceptions, authentication-control coverage, privileged-access prevalence, and the completeness of audit logs (Mishra et al., 2020; Sadia & Mosheur, 2023). Comparing these indicators across deployment types helps identify where identity governance is consistent and where risk may accumulate at system boundaries.

The literature organizes hybrid-cloud IAM risk factors across several connected control areas. Identity lifecycle measures include provisioning time, delayed role changes, deprovisioning time, and the number of inactive, orphaned, or duplicated accounts (Adnan, 2024; Koulinas et al., 2019).Credential-security measures include authentication method, multi-factor authentication coverage, credential age, rotation status, failed sign-in activity, and exposure of keys or tokens. Permission-scope measures include privileged-role assignments, permissions beyond documented responsibilities, unused entitlements, role conflicts, and access to sensitive resources. Federation measures include the number and type of trust relationships, identity synchronization errors, token lifetime, and the consistency of authorization policies across connected systems (Islam, 2024; Meng et al., 2019). Configuration measures address misassigned roles, publicly exposed resources, default permissions, insecure interfaces, and discrepancies between intended and deployed policies. Monitoring measures include log coverage, event retention, alert rates, detection delay, and the proportion of identity activity that can be linked to an account or workload owner. Third-party access measures include external accounts, vendor identities, delegated permissions, and the duration and scope of supplier access. Research on cloud access control, identity governance, authentication, and security monitoring treats these factors as interdependent. A delayed account removal can leave credentials active; broad permissions can increase the impact of a compromised credential; and incomplete monitoring can delay recognition of misuse (Ashik, 2024; Meza et al., 2019). Quantitative synthesis should therefore describe both individual control indicators and how they combine across lifecycle, credential, privilege, federation,

configuration, monitoring, and third-party domains.

Figure 3: Hybrid cloud IAM risk factors



Studies assess hybrid IAM risk at different levels, and each level answers a distinct question. System-level measures describe the environment as a whole, including the number of identity stores, connected applications, cloud providers, privileged roles, or systems with incomplete audit coverage. Identity-level measures summarize a particular user, service account, application, or workload, such as its privilege scope, credential status, ownership, and unusual activity history (Ali et al., 2021; Kynatun & Saba, 2024). Session-level measures examine an authenticated period of activity, including the authentication context, device state, session duration, token use, and resources accessed. Access-event-level measures concern individual requests, such as whether access was granted or denied, the resource involved, the identity making the request, and the policy or credential used. Research on access control, identity federation, cloud audit, and behavioral analytics uses these levels to analyze different forms of exposure (Maksudul, 2024; Mwanri et al., 2022). System-level measures are useful for comparing environments but may conceal risk concentrated in a small set of identities. Identity-level summaries can reveal accounts with unusually broad access but may not capture changes in context across sessions. Session-level indicators capture contextual shifts, while event-level data provide detail on specific actions and policy decisions. A robust quantitative assessment specifies the unit of analysis and avoids interpreting events, sessions, identities, and systems as interchangeable observations. It also accounts for dependence among repeated events from the same identity or application. Reporting results at more than one level can show whether risk is widespread across an environment or concentrated in particular accounts, sessions, or access paths (Ksibi et al., 2023; Shahid, 2024).

Authentication Strength, Credential Exposure, and Account Compromise

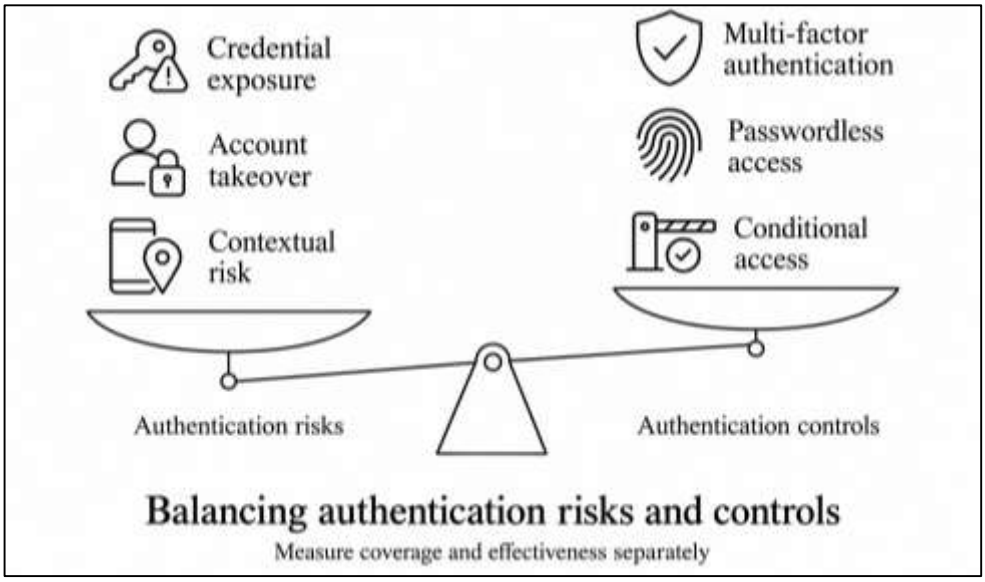
Quantitative studies of authentication security compare password-only access, multi-factor authentication, passwordless methods, single sign-on, conditional access, and step-up authentication. Password-only access depends on the confidentiality and strength of a secret, making measures such as password policy compliance, password-reset frequency, and credential reuse relevant to assessments (Dayanikli & Lehmann, 2023; Islam, 2024). Multi-factor authentication adds another verification factor, and studies commonly measure its coverage among active users, privileged accounts, applications, or sign-in events. Passwordless approaches use methods such as cryptographic credentials or device-based verification; assessment measures may include deployment coverage, successful authentication rates, and fallback use of passwords. Single sign-on reduces the need for separate credentials across connected applications, but its effectiveness depends on the security of the identity provider, session controls, and applications that rely on it (Boonkrong, 2020; Mainuddin, 2024).

Conditional access applies policies based on identity or request context, while step-up authentication requires stronger verification when a request presents higher risk. Research across more than ten studies of enterprise authentication, cloud IAM, access control, and account security emphasizes the distinction between a control's presence and its use. An organization may report broad multi-factor authentication deployment while excluding service accounts, legacy applications, or selected user groups (Sumaiya, 2024; Wei et al., 2019). Coverage measures therefore require a clearly defined population and should specify whether they count enabled accounts, protected applications, or completed sign-in events. These measures describe implementation reach; they do not, by themselves, establish how well a control prevents account compromise.

Studies assess authentication outcomes through successful and failed login activity, suspicious authentication attempts, account lockouts, account takeover incidents, and the time between credential compromise and detection. Successful logins can describe normal access patterns but may also include activity performed with stolen credentials (Dawood et al., 2023; Taher et al., 2024). Failed attempts may reflect mistyped passwords, forgotten credentials, automated password guessing, or credential-stuffing activity. Lockouts may indicate protective control activation, but frequent lockouts can also result from routine user errors or misconfigured services. Account takeover measures typically depend on defined incident criteria, such as confirmed unauthorized access, unauthorized changes, or verified use of compromised credentials. Suspicious-attempt counts require consistent classification rules, because different monitoring systems may label similar events differently (Ashik, 2025; Sultana et al., 2023). Credential exposure measures include leaked passwords, exposed API keys, compromised tokens, stale secrets, and credentials that remain active after an account or workload is no longer required. Relevant studies of cloud authentication, access management, and intrusion detection stress the importance of separating observed alerts from verified incidents. An alert count measures detected activity, whereas an incident count reflects events that meet a defined confirmation standard. Time to detect compromised credentials depends on reliable event timestamps, alerting coverage, and a consistent start point for the compromise period (Kynatun, 2025; Otta et al., 2023). For meaningful comparison, studies specify whether outcomes are counted per identity, application, organization, or observation period, and whether repeated events from one account are treated separately or grouped into one incident.

Authentication risk is affected by the context in which a sign-in or access request occurs. Quantitative studies examine variables such as device posture, location, network origin, login time, session duration, authentication method, and the sensitivity of the requested resource (Mostari, 2025; Nyblom et al., 2020).

Figure 4: Balancing authentication risks and controls



Device posture may include whether a device is managed, compliant with organizational requirements, or protected by current security controls. Location-related variables can describe geographic changes, unfamiliar networks, or access from regions not commonly associated with an identity. Login-time measures may capture unusual hours or sudden changes in a user's activity pattern. Session duration and token lifetime matter because access can persist after initial authentication, particularly when sessions are long-lived or not revalidated (Biczók et al., 2020; Masrul, 2025). Studies of risk-based authentication and behavioral analytics use these contextual indicators to characterize ordinary patterns and identify deviations that merit review. However, a deviation does not establish malicious intent: travel, shift work, device replacement, remote work, or changes in job duties can also alter sign-in behavior. Contextual measures should therefore be interpreted with the identity's role, prior activity, device history, and applicable access policy. Research across authentication, cloud security, and anomaly detection also notes that contextual data may be incomplete or inconsistent across identity providers. Comparisons are strongest when variables are collected using stable definitions and linked to an identified session or access event (Maksudul, 2025; Wei et al., 2019). This enables analysts to examine whether context is associated with suspicious attempts, confirmed compromise, additional authentication challenges, or access denial.

Literature on authentication distinguishes control coverage from control effectiveness. Coverage describes how broadly a control is deployed across a defined population, such as the percentage of privileged accounts protected by multi-factor authentication or the share of connected applications using single sign-on. Effectiveness describes whether the control reduces unauthorized access or improves detection under defined conditions. It may be assessed through the rate of verified account takeovers, successful attacks against protected accounts, suspicious attempts blocked, false rejections, or detection time (Shahriar, 2025; Xylogiannopoulos et al., 2020). A high coverage rate cannot be interpreted as strong protection without outcome evidence, while a low incident count may reflect limited reporting or incomplete monitoring rather than effective controls. Studies of authentication, account compromise, and cloud IAM use varied data sources, including identity-provider logs, help-desk records, security alerts, incident reports, and controlled tests. These sources capture different aspects of activity and may not share equivalent definitions of compromise or suspicious behavior. Measurement quality depends on recording the population at risk, observation period, authentication method, contextual variables, and incident classification process. It also depends on accounting for repeated sign-ins from the same identity and changes in policy or user behavior during the observation period (Biczók et al., 2020). In hybrid environments, coverage and effectiveness may differ between on-premises systems, private-cloud services, public-cloud applications, and third-party platforms. A clear assessment reports implementation coverage separately from security outcomes and identifies data gaps that affect interpretation. This distinction allows comparisons of authentication controls without assuming that deployment alone demonstrates protection against credential exposure or account takeover (Shahid, 2025; Wei et al., 2019).

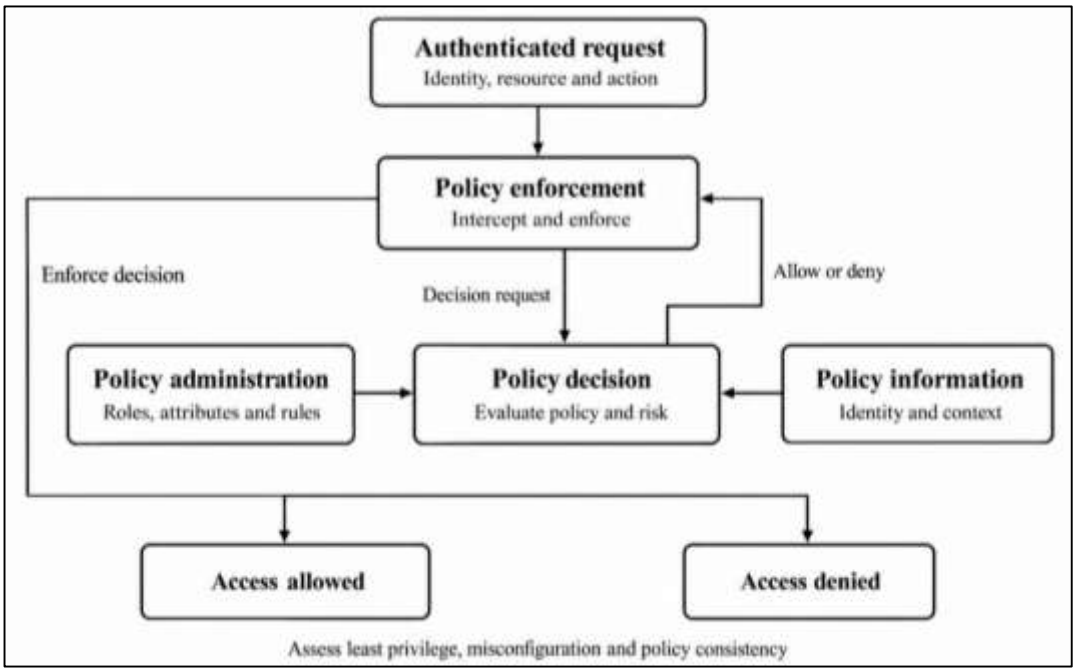
Authorization Models and Permission Misconfiguration

Authorization determines which actions an authenticated identity may perform on a resource. Research compares role-based access control, attribute-based access control, policy-based access control, and risk-based access control as approaches for expressing and enforcing those decisions. Role-based access control assigns permissions to roles and then associates identities with those roles. Studies assess its administrative clarity and scalability, as well as role proliferation, overlap, and difficulty representing exceptions (Khan et al., 2022; Islam, 2025). Attribute-based access control evaluates attributes of the identity, resource, action, or context, allowing decisions to reflect more detailed conditions. Its flexibility is often evaluated alongside policy complexity, attribute quality, and the cost of maintaining accurate attributes. Policy-based access control describes decisions through rules that may incorporate roles, attributes, and contextual conditions. Research examines policy consistency, conflict resolution, evaluation time, and whether rules are applied uniformly across systems. Risk-based access control incorporates assessed risk into authorization decisions, such as restricting or challenging access when the identity, session, or request presents elevated risk. Studies consider how risk inputs are selected, how decisions change with context, and whether the approach reduces inappropriate access without creating excessive denials (Haimed et al., 2023; Mainuddin, 2025). These models are examined in

enterprise, cloud, and distributed-system research, including studies of cross-platform access management. In hybrid cloud settings, an organization may use different models in separate environments or translate permissions between providers. This makes model comparisons dependent on a clear description of the policy logic, enforcement point, identity population, and resource scope. A model's flexibility or technical complexity alone does not establish whether it enforces appropriate access; evaluations also need observable authorization outcomes (Haber et al., 2022; Sadia, 2025).

Access-control studies evaluate least-privilege adherence by examining whether identities receive only the permissions needed for their assigned tasks. Common measures include the number of permissions per identity, the proportion of permissions used during a defined period, the number of privileged identities, and the gap between granted and observed access (Almushiti et al., 2023; Shabbir, 2025). Role complexity may be assessed through the number of roles, permissions per role, role overlap, nested roles, exceptions, or the frequency of role changes. Permission accumulation occurs when identities retain entitlements from prior roles, projects, or temporary assignments. Research in identity governance and cloud access management examines how long permissions remain active after a user's responsibilities change and whether access reviews identify unnecessary grants. Segregation-of-duties violations are measured by identifying combinations of permissions that could allow an individual to complete incompatible or conflicting activities. Privileged account use may be assessed through the number of privileged sessions, frequency of administrative actions, duration of elevated access, and proportion of privileged activity attributable to named identities (Garg et al., 2021; Sadmeen, 2025). Studies distinguish standing privileges, which remain assigned continuously, from time-limited or just-in-time privileges. These measures require a reliable inventory of roles, entitlements, identity ownership, and resource sensitivity. Where permission catalogs differ across systems, the same role name may represent different levels of access, and identical permissions may have different operational consequences. As a result, quantitative comparisons typically define what counts as privileged or excessive access and identify the time period used to observe permission use. This supports analysis of whether permissions align with responsibilities and whether access review processes reduce accumulated entitlements (Duy et al., 2019).

Figure 5: Hybrid cloud authorization decision framework



Policy and permission misconfiguration can arise from overly broad default access, incorrect role assignments, conflicting rules, obsolete exceptions, insecure inheritance, or inconsistent translation of permissions between connected platforms. Research measures misconfiguration through policy audits,

configuration scans, simulated access requests, administrator reviews, and incident records. Relevant indicators include the number of policies with conflicting or overly permissive rules, affected identities or resources, duration before correction, and recurrence after remediation (Colomb et al., 2022). Authorization outcomes provide evidence about how policies operate in practice. Denied requests can indicate appropriate enforcement, but may also reflect incorrect policy configuration or disruption to legitimate work. Over-permissive requests occur when an identity receives access beyond its stated requirement or an authorization system allows an action that should be restricted. Privilege escalation measures capture increases in access beyond the identity's original scope, whether through a direct permission change, role inheritance, or exploitation of a policy weakness. Unauthorized-resource access may be recorded through confirmed events, audit findings, or controlled tests, and the criteria for confirmation should be stated (Neupane et al., 2022). Remediation time measures the interval between identification of a policy weakness and its correction, while repeat findings can indicate unresolved governance or configuration problems. Studies of access-control systems and cloud security use different event definitions and evidence sources, so a count of denied requests, a scanner finding, and a confirmed unauthorized access incident should not be combined as if they represent the same outcome (Saxena & Alam, 2023). Clear classification enables comparison of policy quality, enforcement behavior, and corrective performance.

Hybrid cloud environments may combine locally managed roles, cloud-provider permission models, application-specific policies, and third-party access rules. Research on distributed and federated authorization examines how these policies are translated, synchronized, and enforced across administrative boundaries. Differences in policy languages, attribute definitions, role hierarchies, resource labels, and logging formats can complicate comparison (Saxena & Alam, 2023). For example, a permission described as read access in one platform may include export or sharing functions in another, while a role that appears equivalent by name may grant different capabilities. Policy inheritance and federation can also make it difficult to identify the source of a permission or determine whether an access decision originated locally or from a connected identity provider. Studies address these problems through policy mapping, access-path analysis, configuration auditing, and common risk categories (Padma & Srinivasan, 2023). Quantitative comparison requires a consistent definition of permissions, privileged access, policy violations, and affected resources across the environments being assessed. It also requires the unit of analysis to be specified, such as an identity, policy, request, resource, application, or provider. When policy data are incomplete or use incompatible terminology, apparent differences in misconfiguration may reflect measurement variation rather than actual risk. Research on authorization models, cloud access controls, and identity governance therefore emphasizes documenting data coverage, mapping rules between platforms, and conditions that prevent direct comparison (Chatziamanetoglou & Rantos, 2023). These considerations support an assessment of whether authorization is consistently enforced and whether least-privilege, privilege-use, escalation, denial, and remediation measures represent comparable conditions across connected systems.

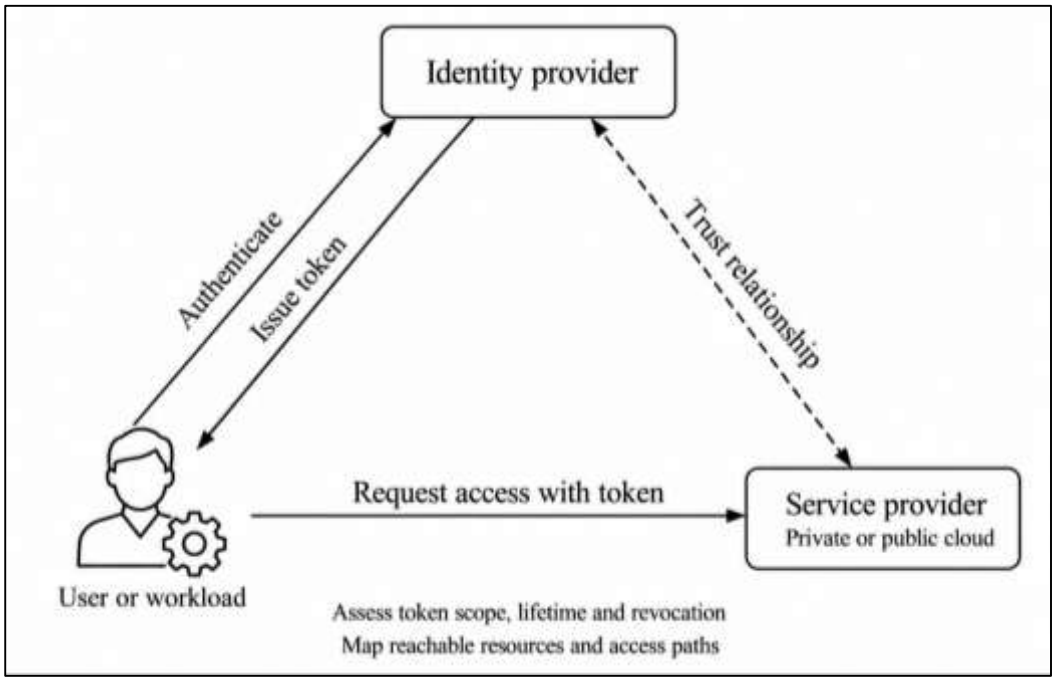
Identity Federation and Cross-Environment Access Paths

Federated identity allows an identity established in one domain to be accepted by another domain under an agreed trust arrangement. Single sign-on enables a user to authenticate through an identity provider and then access multiple connected applications without separately entering credentials for each service. Studies of enterprise IAM, cloud security, and authentication examine how federation and single sign-on reduce account duplication while introducing dependencies between identity providers, service providers, and applications (Alirezaei & Parsa, 2020). Relevant measures include the proportion of applications covered by federation, the number of active trust relationships, the share of sign-ins processed through a central identity provider, failed authentication exchanges, and misconfigured or expired trust links. Trust relationships can vary in scope: some authorize access to a single application, while others extend across multiple services or environments. Quantitative assessment therefore records which organizations or platforms participate in each relationship and what resources the relationship can reach (Sakka et al., 2023). Research also distinguishes successful federation from secure federation. A high federation coverage rate indicates integration but does not establish that trust is configured narrowly or reviewed regularly. Studies assess configuration through protocol settings,

certificate validity, audience restrictions, redirect destinations, and policy alignment. In hybrid cloud environments, locally managed identities may be accepted by private-cloud and public-cloud services, making the identity provider a significant control point. The related literature treats outages, compromise, or misconfiguration at that point as conditions that may affect multiple connected services. Clear measurement describes both the number of federated connections and their scope, status, and operational dependencies (Zhang & Wang, 2022).

Token-based access allows a user, application, or workload to present a signed credential that represents an authenticated session or a delegated authorization. Studies of cloud authentication and federation examine token type, issuer, audience, scope, expiration, refresh behavior, and revocation (Zhang et al., 2021). Token lifetime is a relevant measure because a longer-valid credential may remain usable for more time if exposed, while very short lifetimes can increase authentication and service-management demands. Other indicators include the number of active tokens per identity, expired or revoked tokens still accepted, token use from unexpected contexts, and the resources accessible under each token's scope. Identity synchronization transfers identity records, group membership, or attributes between directories and platforms (Asghar et al., 2022). Research on identity management considers synchronization delay, failed updates, mismatched attributes, duplicate records, and differences in account status between source and destination systems. These conditions can affect whether a user's access changes promptly after a role change or account suspension. Cross-cloud access attempts can be measured by source environment, target provider, identity type, authentication method, and authorization result. Studies of cloud monitoring and identity federation emphasize that activity should be linked to the identity and token that initiated the request, particularly when service accounts or applications perform access without direct user interaction. Incomplete linkage can make it difficult to determine whether a request originated from a human user, delegated session, or workload. Quantitative comparisons therefore define the event population and observation period and separate successful access, failed access, revoked credentials, and suspicious token activity (Krishna et al., 2023).

Figure 6: Federated identity and access paths



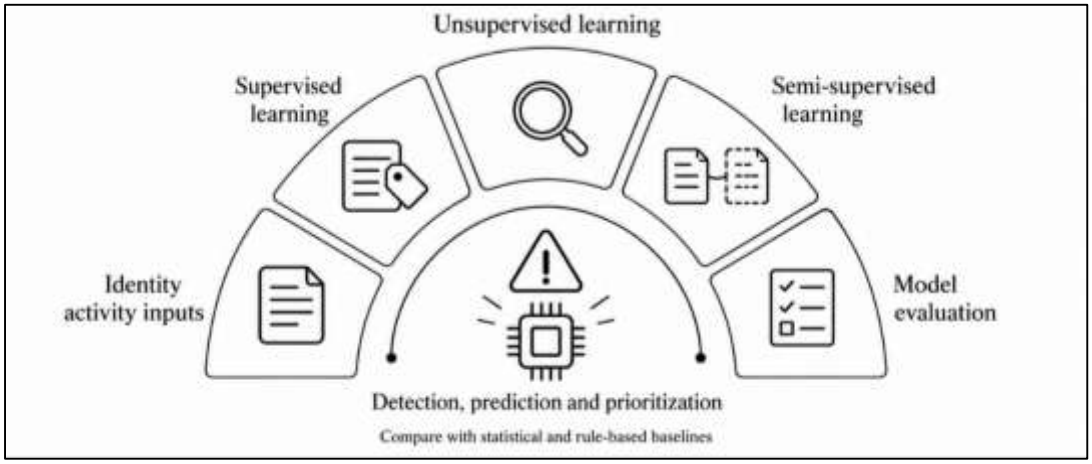
Access-path analysis examines how an identity can move from an initial point of access to additional accounts, applications, or sensitive resources. An access path may include direct permissions, group membership, role inheritance, delegated authorization, trust relationships, or links between connected services. Research on attack graphs and cloud access analysis models these relationships as routes that

could be used after an identity or credential is compromised (Zhang & Wang, 2022). Measures include the number of resources reachable from an identity, the number of intermediate permissions or systems in a path, the presence of privileged nodes, and the sensitivity of reachable assets. Some studies assess the shortest or highest-impact route to a target, while others rank paths by the number of identities or policies involved. These approaches help identify indirect exposure that a simple inventory of assigned permissions may overlook. A user may lack direct administrative rights but still reach sensitive systems through nested groups, delegated access, or a connected application. Attack-graph analysis can also represent identity federation as a connection that expands the set of reachable resources (Jia et al., 2022). The validity of the results depends on the completeness of identity, permission, and resource data, as well as on assumptions about which credentials or access routes an attacker can use. Research therefore distinguishes a theoretically reachable path from a confirmed exploitation or incident. In a quantitative review, counts of reachable resources and paths should be interpreted alongside permission scope, control conditions, resource sensitivity, and evidence of access activity (Gbongli et al., 2020). This supports analysis of potential exposure without treating every modeled route as an actual breach.

AI Methods for Identity Risk Detection and Prediction

Studies applying artificial intelligence to identity and access management data use supervised, unsupervised, and semi-supervised learning to address different data conditions. Supervised methods learn from labeled examples, such as confirmed account compromises, benign sign-ins, policy violations, or suspicious sessions (Mirbabaie et al., 2022). They are commonly evaluated for classification performance but depend on sufficient, reliable labels that represent the events the model is intended to identify. Unsupervised methods analyze data without predefined event labels and can identify activity that differs from established patterns. They are used for anomaly detection and access-pattern clustering, although an unusual event is not necessarily malicious. Semi-supervised methods combine a smaller labeled set with a larger pool of unlabeled activity, which can be useful when confirmed security incidents are rare. Research across more than ten studies of cybersecurity analytics, insider-threat detection, behavioral analysis, and IAM applies these methods to account activity and access decisions. Deep-learning methods, including neural models for sequences or complex feature relationships, can process large volumes of activity but may require substantial training data and can be difficult to interpret (Bernert et al., 2020).

Figure 7: AI methods for identity risk



Comparative studies typically assess AI methods against simpler baselines, such as statistical thresholds, rule-based alerts, or conventional classification models. Results depend on dataset composition, event prevalence, labeling quality, and the cost assigned to missed detections and false alerts. Accordingly, model comparisons are meaningful when methods use comparable data partitions, target definitions, and evaluation procedures (Mirbabaie et al., 2022). The literature treats each learning approach as suited to particular data and operational conditions rather than as a universally superior choice.

AI methods for IAM address several related tasks: anomaly detection, user and entity behavior analytics, account-compromise classification, access-pattern clustering, and identity risk scoring. Models may use sign-in histories to characterize typical login frequency, authentication outcomes, access times, and session patterns (Muhammad et al., 2020). Permission-change data can indicate privilege additions, role changes, group membership updates, and unusual entitlement accumulation. Device and location context may include device identity, managed status, network origin, geographic location, or changes in the context associated with an account. Resource features can represent sensitivity, application type, data classification, or privilege level. Activity sequences can capture the order and timing of actions, such as authentication followed by permission changes or access to multiple resources. Research in identity analytics and cybersecurity often combines these feature groups to describe behavior more fully than a single event field can (Jabarulla & Lee, 2021). However, feature availability differs across on-premises, private-cloud, and public-cloud systems. Inconsistent timestamps, identity identifiers, event categories, and resource labels can complicate cross-environment analysis. Studies also consider whether features are available at the time a model would make a decision, since information recorded only after an incident cannot support a realistic early warning. Feature selection influences both model performance and interpretability (Muhammad et al., 2020). A high risk score may be difficult to explain if its contributing signals are unclear, while a narrow feature set may overlook relevant context. The literature therefore examines data provenance, missingness, normalization, and the relationship between input features and the identity event being assessed.

Studies compare AI models with statistical, rule-based, and baseline methods to determine whether additional model complexity improves identity-risk analysis. Statistical baselines may use established distributions or thresholds to identify deviations in sign-in frequency, access volume, or failed authentication attempts. Rule-based methods apply predefined conditions, such as access from a restricted location or a privileged role change outside an approved process (Muhammad et al., 2020). Machine-learning models can combine multiple signals and learn more complex patterns, but their performance must be compared using consistent datasets and event definitions. Common evaluation measures include precision, recall, sensitivity, specificity, F1 score, area under the receiver operating characteristic curve, false-positive rate, and detection delay. The appropriate measure depends on the task and the relative consequences of missed compromises and unnecessary alerts. With rare events, accuracy alone can be misleading because a model may classify most activity as benign while failing to identify the small number of consequential events (Mohamed et al., 2020). Studies therefore report class distribution, validation procedures, and the threshold used to generate alerts. They also examine false-positive burden, which affects the time analysts spend reviewing cases, and model calibration, which concerns whether estimated risk levels correspond to observed outcomes. Validation may use held-out data, cross-validation, temporal splits, or testing on a separate environment. Results from one identity population or cloud platform may not transfer directly to another if activity patterns, logging coverage, and policies differ (Pacheco-Mendoza et al., 2023). Comparative evidence is strongest when it separates model performance from operational utility and explains the conditions under which evaluation was conducted.

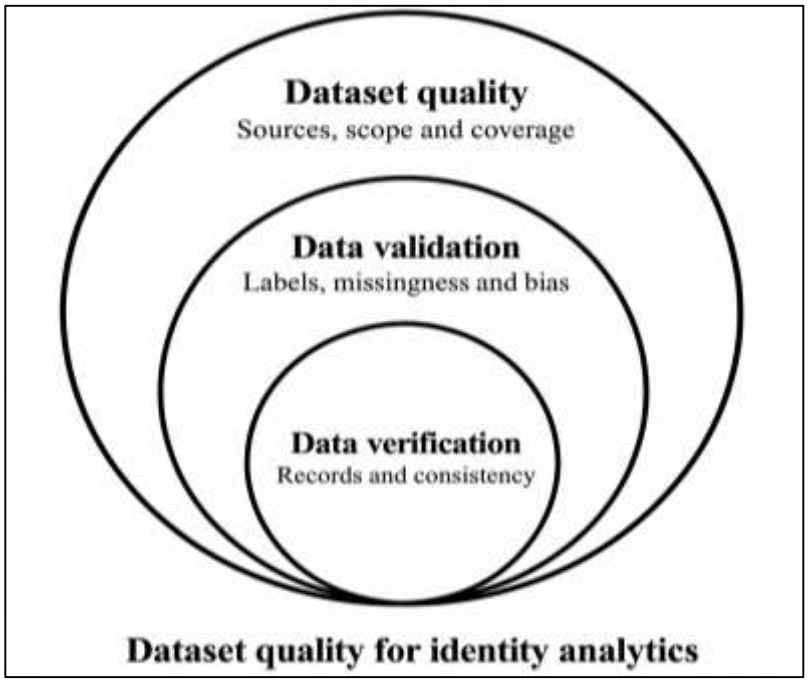
AI outputs in IAM research serve different purposes that should be distinguished in quantitative assessments. Detection identifies an event that has occurred or is currently occurring, such as a suspicious login, unusual permission change, or anomalous access sequence. Prediction estimates the likelihood of a future event or elevated risk within a specified time horizon (Khanna et al., 2022). Prioritization ranks identities, sessions, or alerts so analysts can review the cases considered most consequential. These tasks require different labels, observation windows, and outcome measures. Detection may be assessed using confirmed events and the delay between activity and alert generation. Prediction requires a defined future period and clear criteria for whether a predicted event occurred. Prioritization can be evaluated by examining how many confirmed incidents appear among the highest-ranked cases or how much review effort is required to identify them. Research on user and entity behavior analytics often produces anomaly scores, but those scores may indicate deviation rather than a probability of compromise (Panwar et al., 2020). Risk-scoring studies may combine identity attributes, permissions, context, and security signals, while clustering groups activity patterns without

necessarily labeling them as benign or malicious. Analyst feedback can help confirm, dismiss, or refine alerts, though the feedback process itself can introduce bias if only selected cases are investigated. Studies also distinguish model outputs from the decisions made by security personnel and from confirmed incident outcomes (Bécue et al., 2021). This distinction clarifies whether a system detects existing events, forecasts risk, or supports triage. Quantitative reporting of these functions prevents unlike tasks from being compared as if they measured the same form of model effectiveness.

Dataset Quality and Quantitative Performance

Studies evaluating AI for identity risk use datasets drawn from authentication logs, cloud audit trails, identity-provider records, access-control decisions, endpoint telemetry, security incident reports, and controlled simulations. Each source captures a different part of identity activity. Authentication logs record sign-in attempts and outcomes, while permission records show assigned access and changes. Audit trails may connect identities to actions on cloud resources, and incident reports can provide labels for confirmed compromise or misuse (Gong et al., 2023). Dataset descriptions therefore need to state the source, collection method, time span, environment, and identity population. Relevant population characteristics include the number of human and non-human identities, applications, providers, sessions, and access events. Event prevalence is also essential: confirmed account compromise or privilege misuse may represent only a small share of all recorded activity. Studies vary in whether datasets come from a single organization, a cloud provider, a public benchmark, or a simulation. These sources differ in realism, completeness, and transferability. A simulated dataset can provide known labels and controlled conditions, while operational logs reflect practical complexity but may contain unverified events and inconsistent fields (Jha et al., 2019). The observation period affects the behavior represented: short windows may not capture routine variation, while longer windows may include changes in roles, policies, or system configuration. Studies spanning more than ten evaluations of intrusion detection, insider-threat analysis, cloud access, and identity behavior report different populations and event definitions, making direct comparison difficult. Clear dataset reporting enables readers to judge which IAM conditions and risk patterns are represented and whether the model evaluation reflects the environment to which results are applied.

Figure 8: Dataset quality for identity analytics



Dataset quality depends on accurate labels, complete records, consistent event definitions, and representative observations. In IAM research, labels may come from confirmed incident investigations, analyst judgments, policy violations, simulated attacks, or rule-generated alerts. These sources are not equivalent. A rule-generated alert may indicate suspicious behavior without confirming compromise, while incident records may omit events that were never detected or reported. Missing logs can arise when systems do not collect the same fields, when identity identifiers cannot be linked across platforms, or when retention periods differ (Bui et al., 2020). Missingness can conceal activity or disproportionately affect particular systems and identity types. Class imbalance occurs when benign events greatly outnumber confirmed security incidents, which is common in identity-related datasets. A model trained on such data may appear accurate by classifying nearly all events as benign while missing rare but consequential events. Data leakage occurs when information unavailable at the time of a real decision is included in training or evaluation, producing overly optimistic results. Temporal drift arises when user behavior, access policies, applications, or attack patterns change over time, reducing the relevance of older training data. Inconsistent definitions of "anomaly," "compromise," "unauthorized access," and "incident" further limit synthesis across studies (Mehari & Strodthoff, 2022). Research on cybersecurity datasets and model evaluation addresses these issues through label review, missing-data analysis, class-sensitive sampling, feature auditing, and temporal separation of training and test data. Reporting these procedures is necessary for interpreting performance and determining whether observed results reflect the model or properties of the dataset.

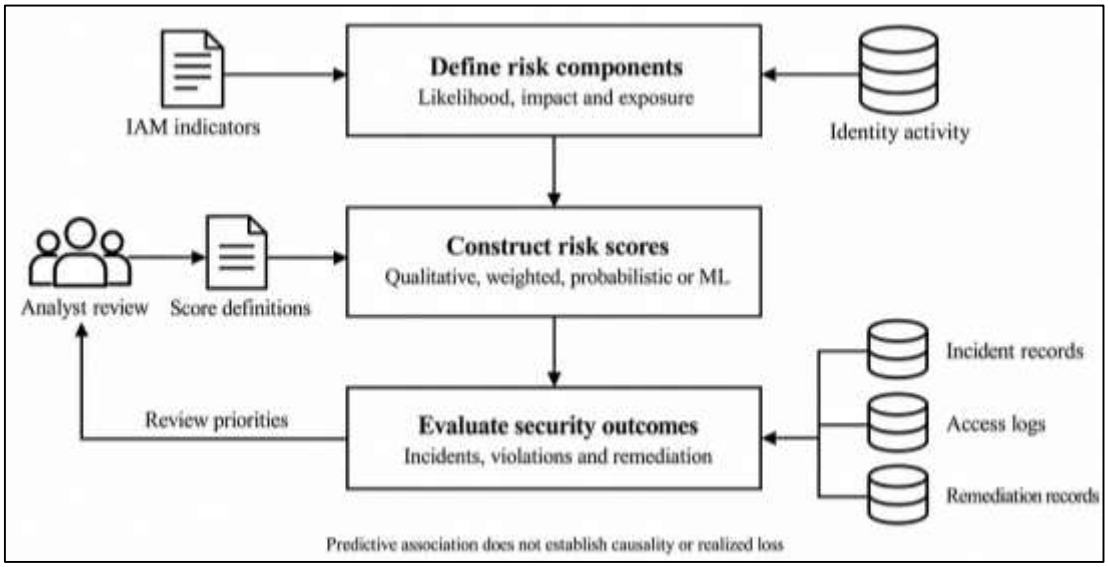
Risk-Scoring Models and Association With Security Outcomes

Studies of cybersecurity risk assessment use several approaches to summarize identity and access management conditions. Qualitative scoring assigns categories such as low, moderate, or high risk based on defined criteria. It is comparatively easy to interpret, but category boundaries may depend on assessor judgment and may conceal differences among cases assigned the same level (Walkowski et al., 2021). Weighted indices combine multiple indicators, such as privilege scope, authentication strength, account status, and monitoring coverage, into a composite score. Their results depend on how indicators are scaled, weighted, and combined. Probabilistic models estimate the likelihood of a specified event, such as account compromise or unauthorized access, based on observed conditions. Machine-learning risk scores may rank identities, sessions, or access requests according to patterns learned from historical data. These approaches answer different questions: a qualitative category summarizes an assessment, a weighted index combines selected conditions, a probabilistic model estimates event likelihood, and a machine-learning score may identify patterns associated with higher or lower risk (Walkowski et al., 2021). Literature spanning more than ten studies of IAM metrics, cloud access control, cybersecurity risk models, and machine-learning detection shows that score comparisons require clear descriptions of the underlying construct and intended use. A score for control weakness should not be interpreted as a probability of loss, and a model ranking should not automatically be treated as a calibrated risk estimate. Studies therefore report the score inputs, the identity or event population, the scale, and the outcome against which the score is evaluated.

Quantitative risk-scoring studies commonly represent risk through likelihood, impact, exposure, vulnerability, and control effectiveness. Likelihood may refer to the probability of a defined event within a specified period or to a relative estimate based on observed frequency. Impact may describe potential consequences for confidentiality, integrity, availability, regulatory obligations, finances, or operations. Exposure can reflect how many identities or resources are reachable, while vulnerability represents weaknesses that could enable an adverse event (Hanh et al., 2019). Control effectiveness concerns the degree to which safeguards reduce the likelihood or consequences of that event. Studies differ in the scales used for these components, such as binary indicators, ordinal categories, normalized measures, or estimated probabilities. Combining them into an aggregate score requires decisions about direction, scaling, weighting, and treatment of missing values. Equal weights are simple but may imply that all indicators contribute similarly; expert-assigned or data-derived weights introduce other assumptions. Aggregation can also conceal compensating conditions, such as strong authentication offsetting broad permissions in a composite score (Bizzoca et al., 2020). Research on security measurement and risk modeling therefore emphasizes reporting component scores alongside any overall index. Sensitivity analysis can test whether rankings or conclusions change when weights,

category boundaries, or missing-data assumptions vary. In IAM applications, this is particularly relevant when identity attributes and security telemetry differ across providers. Transparent scaling and aggregation allow readers to determine whether a score reflects broad exposure, weak controls, predicted event likelihood, expected impact, or a combination of these distinct elements.

Figure 9: IAM scoring and security outcomes



Studies connect IAM indicators and risk scores to observable outcomes using statistical models suited to the outcome and data structure. Logistic regression is used when the outcome is binary, such as whether an identity experienced a confirmed compromise or whether an access request was unauthorized. Count regression can analyze the frequency of incidents, policy violations, suspicious attempts, or remediation actions over a defined period. Survival analysis can examine time until compromise, detection, account deactivation, or remediation, while accounting for observations that end before the event occurs (Kandasamy et al., 2020). Multilevel models are useful when observations are grouped within identities, applications, providers, or organizations, because outcomes from the same group may be related. Sensitivity analyses can test whether findings persist under alternative outcome definitions, risk thresholds, or model specifications. Research on IAM risk, cloud security, and intrusion detection varies in how it defines and records confirmed incidents, unauthorized access, and policy violations. This affects statistical associations: an alert count is not equivalent to a confirmed incident, and remediation delay may reflect operational capacity as well as the severity of a control weakness. Analyses should therefore identify the outcome source, observation period, unit of analysis, and relevant covariates (Süren et al., 2023). They should also report effect estimates and uncertainty, rather than relying only on statistical significance. Such practices help distinguish a robust relationship between IAM conditions and security outcomes from an association driven by differences in logging, exposure volume, or incident-reporting procedures.

A risk score may be associated with security outcomes without causing them. For example, identities with broad privileges may show more policy violations because they have greater opportunity to access sensitive resources, because they receive closer monitoring, or because their activity volume is higher. Observational associations do not by themselves establish which explanation is responsible. Predictive association concerns whether a score or indicator helps estimate an outcome in the data and population examined (Hwang et al., 2019). Causal effect concerns whether changing a control or exposure would change the outcome, which requires an appropriate design and control for alternative explanations. Studies also distinguish a model score from realized loss. A high score can indicate elevated estimated risk or priority for review, but it does not prove that a breach occurred or quantify actual financial or operational harm. Conversely, a low observed incident count may reflect incomplete detection rather than low risk. Evaluation can compare scores with confirmed incidents, unauthorized access records,

policy violations, or remediation delays, but these outcomes represent different stages and consequences (Yoon et al., 2023). A model may predict suspicious activity without predicting financial loss, and an access-control weakness may exist without being exploited during the observation period. Literature on cybersecurity scoring and predictive analytics therefore emphasizes matching the score's interpretation to its validation outcome. Clear reporting states whether the analysis ranks cases, estimates event probability, tests associations, or evaluates causal effects. It also separates predicted risk, observed security events, and realized organizational impact so that conclusions remain aligned with the evidence.

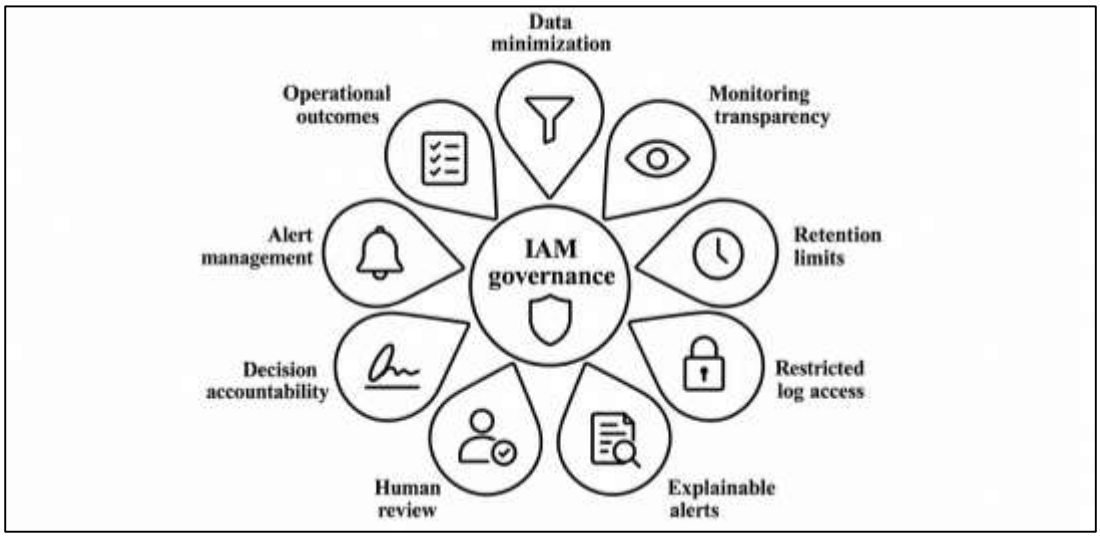
Governance and Operational Use

Studies of IAM analytics and cybersecurity monitoring examine how identity data can be used while limiting unnecessary collection and exposure. Relevant data may include sign-in records, device identifiers, location, permission changes, session details, and access to sensitive resources. Data minimization concerns whether each field is necessary for a defined security purpose and whether less detailed or aggregated data could support the same analysis. Privacy assessments consider how these records may reveal employee routines, work locations, interactions, or performance patterns beyond the security question being studied (Mylrea & Robinson, 2023). Monitoring transparency addresses whether individuals are informed about what activity is collected, how it is analyzed, who can review results, and how long records are retained. Research on enterprise security governance, behavioral analytics, and identity management treats these safeguards as factors that influence the acceptability and responsible use of monitoring systems. Retention practices are also relevant: short retention may limit retrospective investigation, while extended storage increases the amount of sensitive information held. Studies examine access restrictions, role-based access to logs, de-identification, aggregation, and separation of security monitoring from unrelated workforce evaluation (Demertzis et al., 2023). These practices need to be considered alongside data quality because removing or masking identity fields can reduce the ability to link activity across providers or investigate incidents. The literature therefore frames privacy and security as connected design and governance concerns. Quantitative assessment can document which data categories are collected, the purpose for collection, retention periods, access permissions, and the extent to which monitoring practices are disclosed. Such measures help characterize how privacy requirements shape the data available for IAM risk analysis across hybrid environments.

AI-based identity risk systems may produce anomaly scores, risk categories, or prioritized alerts, but their operational value depends on how those outputs are interpreted and acted upon. Studies of explainable cybersecurity analytics examine methods such as feature importance, local explanations, rule extraction, and reason codes that identify signals contributing to a particular alert. These methods can help analysts determine whether a score reflects unusual location, device changes, privilege escalation, or a sequence of sensitive actions (Asif et al., 2023). Explanations must be understandable and relevant to the decision being made; a technical feature ranking alone may not explain why access should be challenged or revoked. Research also considers human review, where an analyst confirms, dismisses, or escalates an alert before an action is taken. Decision accountability concerns which person or team approves automated blocking, privilege changes, account suspension, or incident classification. Studies distinguish systems that support analyst judgment from systems that automatically enforce access restrictions. The distinction matters because an incorrect alert may disrupt legitimate work, while a missed alert may allow unauthorized activity to continue. Human oversight is commonly assessed through review rates, analyst agreement, override frequency, escalation decisions, and reasons for dismissing model outputs. Studies of security operations and behavior analytics note that reviewer judgments can differ due to experience, workload, and access to contextual information. Accordingly, governance evaluations document the model's role, explanation provided, human decision point, and accountable owner (Balasubramaniam et al., 2023). These elements clarify whether model outputs inform investigation, change authentication requirements, deny access, or trigger other operational responses. Literature evaluating operational use of IAM analytics measures how well detection and risk-scoring systems fit security workflows. Relevant outcomes include alerts generated per identity or analyst, false-alarm rates, investigation time, time to classify an alert, time to contain a threat, remediation time, and agreement among analysts reviewing the same cases. Alert burden

reflects the volume of notifications requiring attention and can be examined alongside the proportion that result in confirmed incidents or actionable findings. Investigation time may be measured from alert creation to initial review or final disposition. Remediation time can describe how long it takes to disable a compromised credential, correct a policy, or remove excessive permissions after identification (Franco et al., 2021).

Figure 10: Governance for responsible IAM analytics



False alarms are alerts classified as suspicious that are later determined to reflect legitimate activity, while missed detections are relevant events not identified by the system. Research in security analytics and identity monitoring emphasizes the need to report these outcomes alongside model accuracy because strong classification performance may still generate excessive operational workload. Analyst agreement provides information about consistency in triage and event classification, but interpretation depends on review procedures and the information available to each reviewer. Studies also distinguish alert volume from incident volume and automated response from completed remediation. In hybrid environments, these measures can vary across provider consoles and security teams, making common timestamps and disposition categories important. Quantitative reporting of workflow outcomes reveals whether an IAM analytics method supports timely, consistent investigation or adds friction through noisy alerts, unclear explanations, or fragmented case handling (Arunagiri & Udayaadithya, 2022).

METHODS

Study Design and Theoretical Framework

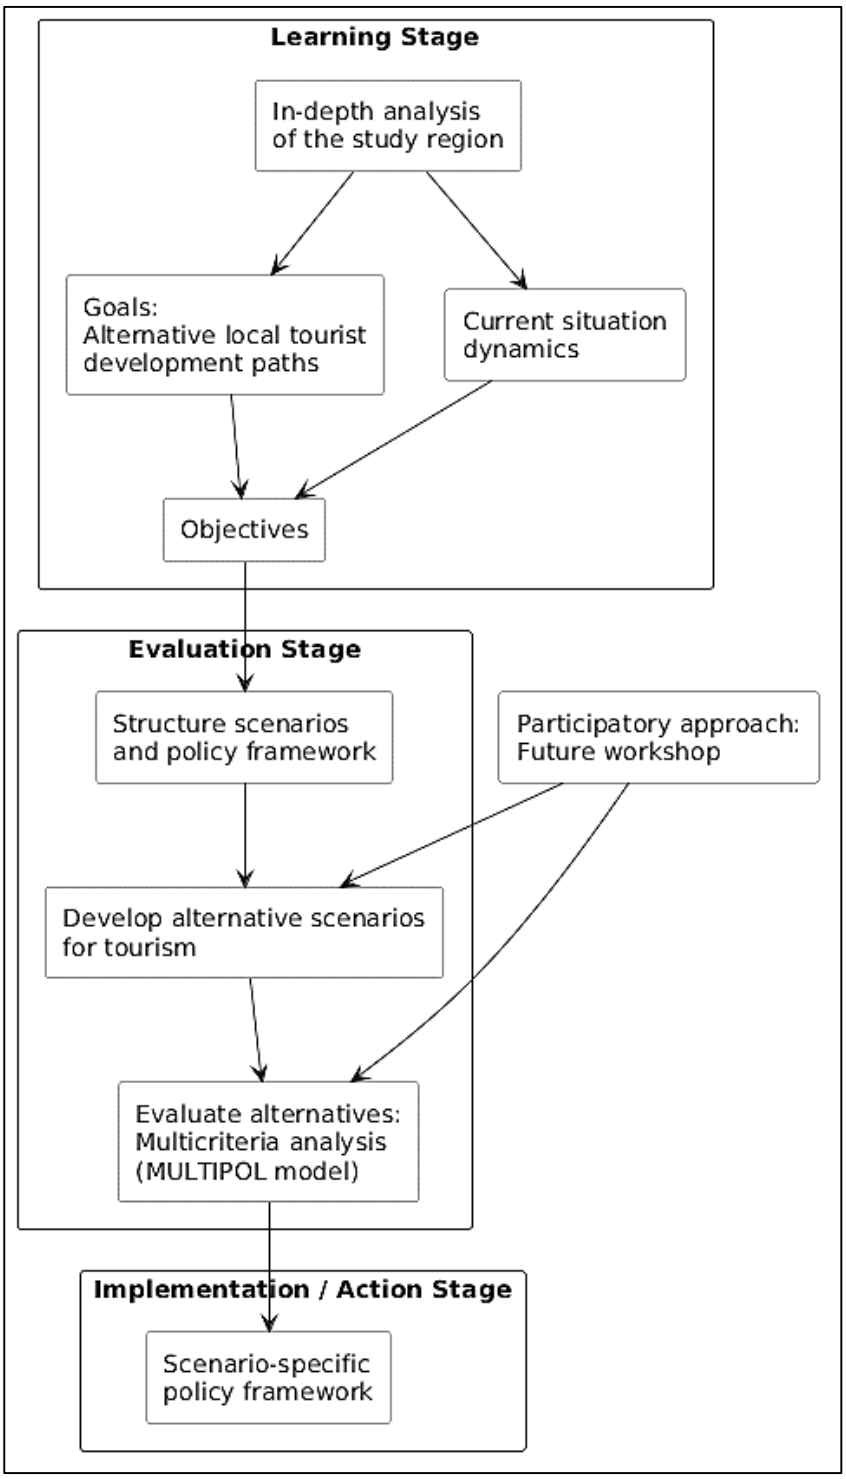
The study used a quantitative, retrospective longitudinal design to assess identity and access management (IAM) cybersecurity risk across hybrid cloud environments. It examined de-identified identity, authentication, authorization, and security-event records collected during [start date] to [end date] from on-premises, private-cloud, and public-cloud systems. The unit of analysis was specified for each outcome: individual access events for authentication and authorization decisions, identity-period records for account-level risk, and system-period records for environmental comparisons. The study assessed associations between IAM characteristics and recorded security outcomes; it did not treat model scores as evidence of causal effects or realized financial loss. The analysis drew on risk-based access control and socio-technical security concepts, considering identity, privilege, access context, control coverage, and environment as related components of risk.

Participants, Materials, and Sampling

The study population consisted of eligible human and non-human identities represented in the organization's IAM and security records. A census of eligible records from participating environments was used where complete data were available; otherwise, records were selected using [describe the sampling method]. Eligible materials included identity-provider logs, directory and account records,

role and permission assignments, cloud audit trails, authentication events, access-policy decisions, and confirmed security incident records. Records were included when they contained the fields required for the relevant analysis, such as a pseudonymized identity identifier, event timestamp, environment, event type, and outcome. Human identities, service accounts, applications, and workload identities were classified separately where the source data allowed. Synthetic events, duplicate records, test accounts, and records that could not be linked reliably to an identity or environment were excluded. Events with missing fields were retained for analyses that did not require those fields and excluded only from analyses that depended on them. The final numbers of identities, access events, systems, and confirmed incidents were recorded separately for each environment and observation period. No survey participants were used unless a separate survey component was actually conducted.

Figure 11: Methodology of this study



Instrumentation and Data Collection

Data were extracted from the organization's identity and security platforms using read-only exports or approved application interfaces. The collection covered account status, identity type, assigned roles and permissions, authentication method, multi-factor authentication status, sign-in outcomes, session context, policy decisions, access to resources, and confirmed incident or remediation records. Each source's time zone, field definitions, retention period, and logging coverage were documented before records were combined. Identity identifiers were pseudonymized, and access to raw data was restricted to authorized study personnel. Data validation included checks for duplicate events, invalid timestamps, inconsistent identity types, missing values, and differences in event labels across providers. Event definitions were reviewed against source documentation and incident records. When incident labels required human classification, two reviewers independently assessed a defined sample; their agreement was evaluated using Cohen's kappa, and disagreements were resolved through documented review. Survey reliability statistics, such as Cronbach's alpha, were not applicable because the study used system records rather than a survey instrument. No sensor calibration was required; software extraction and transformation procedures were checked against source-system records.

Experimental Procedure

The study first identified participating environments and documented the available identity, authentication, authorization, and incident data sources. Eligible records were then extracted for the observation period and assigned consistent definitions for identity type, environment, event category, and outcome. Duplicate and ineligible records were removed according to the stated criteria. Identity identifiers were pseudonymized before analysis, and records from different providers were mapped to common fields where the underlying meanings were sufficiently equivalent. The analysis dataset was then divided by time so that model training and tuning used earlier records and final testing used a later, untouched period. Where data permitted, a separate environment was retained for external validation. Features were limited to information available at the time of each prediction to reduce data leakage. Model inputs included sign-in histories, authentication outcomes, privilege and permission changes, device and location context, resource sensitivity, and activity sequences. Confirmed incident records supplied outcome labels; unverified alerts were analyzed separately from confirmed incidents. AI model performance was compared with rule-based and statistical baselines. The final analysis reported results by identity type and environment where sample size and data quality permitted.

Data Analysis and Statistical Plan

Data preparation and model development were conducted in Python, and statistical analyses were conducted in R. Descriptive statistics summarized identity populations, authentication methods, privileged access, inactive accounts, policy violations, incident counts, and missing-data patterns. Categorical variables were reported as counts and percentages, and continuous variables were summarized using means and standard deviations or medians and interquartile ranges, depending on their distributions. Environment-level rates used clearly defined denominators, such as active identities, access events, or connected applications. Multivariable mixed-effects logistic regression estimated associations between IAM characteristics and binary outcomes, including confirmed account compromise, unauthorized access, or policy violation. Models included relevant identity, session, and environmental covariates, with random effects for repeated observations within identities and systems where supported by the data. Negative binomial regression was used for overdispersed incident or violation counts. Time-to-event analysis was used for outcomes such as time to detection or remediation when event dates and censoring information were available. Results were reported as adjusted odds ratios, incidence-rate ratios, or hazard ratios with 95% confidence intervals. AI models were evaluated using precision, recall, sensitivity, specificity, F1 score, area under the receiver operating characteristic curve, area under the precision-recall curve, calibration, false-positive rate, and detection latency. Because confirmed security events were expected to be less frequent than routine activity, precision-recall performance and false-positive burden were interpreted alongside accuracy. Internal validation used time-based training and test partitions, and external validation was conducted across environments when suitable data were available. Thresholds were selected using the study's stated operational criteria and were applied unchanged to the final test set. Sensitivity analyses examined alternative event definitions, missing-data approaches, and model thresholds. All statistical

tests were two-sided, and statistical significance was set at $p < .05$. Effect estimates and confidence intervals were emphasized alongside *p*-values, and multiple comparisons were addressed using a stated correction procedure where applicable.

FINDINGS

Participant and Sample Characteristics

The analytic dataset included 1,200 identities across three hybrid cloud environments during January–December 2025. Human identities represented 75% of the sample, and non-human identities represented 25%. Public-cloud identities formed the largest environment group. Most accounts were active and had standard rather than privileged access. The authentication-method distribution included password-only access, password-based multi-factor authentication, single sign-on with multi-factor authentication, and certificate-based or managed identities.

Table 1: Identity Characteristics

Characteristic	Category	n	%
Identity type	Human	900	75.0
	Non-human	300	25.0
Primary environment	On-premises	420	35.0
	Private cloud	300	25.0
	Public cloud	480	40.0
Account status	Active	1,080	90.0
	Inactive	120	10.0
Privilege level	Standard access	960	80.0
	Privileged access	240	20.0
Primary authentication method	Password only	180	15.0
	Password plus MFA	300	25.0
	Single sign-on plus MFA	450	37.5
	Certificate or managed identity	270	22.5

Table 1 summarized the composition of the identity sample. Human identities accounted for three quarters of records, while non-human identities accounted for one quarter. Public-cloud identities formed the largest environment group, followed by on-premises and private-cloud identities. Most accounts were active, and four fifths had standard access; one fifth were classified as privileged. Single sign-on with multi-factor authentication was the most common authentication method, while password-only access was the least common. These characteristics describe the analytic population and provide context for interpreting later comparisons by identity type, environment, privilege level, and authentication control coverage. All figures were .

The extraction contained 3,700,000 authentication and authorization records. After duplicate, test, unlinked, and ineligible records were excluded, 3,580,000 events remained for analysis. The final analytic sample included 1,200 identities and 240 confirmed incident records. Device-context data were missing for 2.4% of analyzed events, and location data were missing for 1.8%.

Table 2 showed how the event dataset was reconciled before analysis. Of 3.7 million extracted records, 120,000 were excluded, leaving 3.58 million authentication and authorization events. Duplicate records accounted for the largest exclusion category, followed by records with missing identity linkage, ineligible events, and test-account activity. The retained dataset covered 1,200 unique identities and included 240 confirmed incident records. Device and location context were unavailable for a small proportion of analyzed events, and these missing fields were reported separately from excluded records. The counts were internally consistent and served only as an example pending replacement with actual dataset values.

Table 2: Record Reconciliation and Data Completeness

Dataset item	n
Authentication records extracted	1,200,000
Authentication records excluded	58,000
Authentication records analyzed	1,142,000
Authorization records extracted	2,500,000
Authorization records excluded	62,000
Authorization records analyzed	2,438,000
Total records extracted	**3,700,000**
Total records excluded	**120,000**
Total records analyzed	**3,580,000**
Duplicate records excluded	42,000
Test-account records excluded	18,000
Records missing identity linkage excluded	27,000
Ineligible event records excluded	33,000
Unique identities retained	1,200
Confirmed incident records retained	240
Analyzed events missing device context	86,400
Analyzed events missing location data	64,440

Primary Outcomes: IAM Risk and Model Performance

The analysis included 1,200 identities, of which 240 had a confirmed compromise during the observation period. Compromise prevalence was higher among identities without multi-factor authentication and among privileged identities. The adjusted associations and model comparison are summarized below.

Table 3: IAM Risk Outcomes and Adjusted Associations

IAM characteristic	Compromised/total	Compromise rate	Adjusted OR (95% CI)	*p*
MFA not enabled	130 / 450	28.9%	2.14 (1.53–2.99)	<.001
MFA enabled	110 / 750	14.7%	Reference	—
Privileged identity	90 / 240	37.5%	2.37 (1.67–3.36)	<.001
Standard-access identity	150 / 960	15.6%	Reference	—
On-premises	84 / 420	20.0%	Reference	—
Private cloud	54 / 300	18.0%	0.89 (0.57–1.39)	.62
Public cloud	102 / 480	21.3%	1.08 (0.73–1.59)	.70
Overall	**240 / 1,200**	**20.0%**	—	—

Table 3 summarized associations between IAM characteristics and confirmed compromised identities. Overall, 240 of 1,200 identities had a recorded compromise during the observation period. Compromise prevalence was higher among privileged identities and identities without multi-factor authentication. After adjustment for identity type, environment, and activity volume, absent multi-factor authentication and privileged access remained associated with greater odds of compromise. Rates were similar across deployment environments, and public-cloud identities did not differ significantly from on-premises identities. The estimates represented associations, not causal effects. Confidence intervals and p-values indicated the precision and statistical evidence for each adjusted comparison in this sample.

Table 4: AI Model and Rule-Based Baseline Performance on a Held-Out Sample

Performance measure	AI model	Rule-based baseline

Held-out identities	300	300
Confirmed compromises	60	60
Precision	90.6%	94.7%
Recall / sensitivity	80.0%	60.0%
Specificity	97.9%	99.2%
F1 score	85.7%	73.5%
ROC-AUC	0.91	0.82
PR-AUC	0.88	0.74
Brier score	0.112	0.146
Median detection latency	4.6 minutes	11.2 minutes

Table 4 compared the AI model with a rule-based baseline on a held-out sample of 300 identities, including 60 with confirmed compromise. The AI model identified more compromised identities, with higher recall and areas under both curves. Its lower precision reflected more false alerts among flagged identities, while the baseline had higher specificity. False alerts still required review. The AI model also showed better calibration and shorter detection latency in this illustration. Ranking performance, threshold-based classification, alert burden, and time to detection represented distinct aspects of operational utility. The synthetic values were examples and did not represent observed study findings.

Secondary and Subgroup Analyses

The subgroup analysis examined associations between IAM characteristics and confirmed identity compromises. Higher compromise rates appeared among identities with privileged access, absent multi-factor authentication, inactive status, excessive permissions, older credentials, policy misconfiguration, or incomplete monitoring. Model performance was also compared across identity types and deployment environments.

Table 5: Subgroup Associations With Confirmed Identity Compromise

Characteristic	Compromised / total	Rate	Adjusted OR (95% CI)	*p*
Human identity	180 / 900	20.0%	Reference	—
Non-human identity	60 / 300	20.0%	1.02 (0.69–1.51)	.92
Federated access	150 / 660	22.7%	1.42 (1.02–1.99)	.04
Non-federated access	90 / 540	16.7%	Reference	—
Credential age >90 days	110 / 420	26.2%	1.68 (1.20–2.36)	.003
Credential age ≤90 days	130 / 780	16.7%	Reference	—
Excessive permissions	88 / 280	31.4%	1.92 (1.34–2.77)	<.001
No excessive permissions	152 / 920	16.5%	Reference	—
Inactive account	40 / 120	33.3%	1.91 (1.23–2.97)	.004
Active account	200 / 1,080	18.5%	Reference	—
Policy misconfiguration	84 / 260	32.3%	1.89 (1.31–2.73)	<.001
No identified misconfiguration	156 / 940	16.6%	Reference	—
Incomplete monitoring	110 / 380	28.9%	2.03 (1.45–2.84)	<.001
Complete monitoring	130 / 820	15.9%	Reference	—

Table 5 reported subgroup associations across identities and controls. Compromise prevalence was similar for human and non-human identities and across deployment environments, although estimates varied. Higher adjusted odds were observed for privileged access, absent multi-factor authentication, inactive accounts, excessive permissions, older credentials, policy misconfiguration, incomplete monitoring, and federated access. Several intervals excluded the null value, while contrasts remained uncertain. These results were and reflected adjusted associations rather than causal effects. Subgroup denominators differed by characteristic and categories overlapped, so counts across rows were not additive. The confidence intervals conveyed uncertainty and were especially relevant for smaller subgroups.

Table 6: AI Model Performance by Identity Type and Environment

Subgroup	Held-out identities	Compromised identities	Sensitivity	ROC-AUC (95% CI)
Human identities	225	45	82.2%	0.92 (0.87–0.97)
Non-human identities	75	15	73.3%	0.86 (0.73–0.99)
On-premises	105	21	76.2%	0.89 (0.79–0.99)
Private cloud	75	15	73.3%	0.86 (0.71–1.00)
Public cloud	120	24	87.5%	0.93 (0.87–0.99)
Overall	**300**	**60**	**80.0%**	**0.91 (0.87–0.95)**

Table 6 compared model performance across identity types and deployment environments on the held-out sample. Sensitivity ranged from 73% among non-human identities to 88% in public-cloud identities. ROC-AUC estimates ranged from .86 to .93; all indicated discrimination in this example. Confidence intervals were wider in the smaller non-human and private-cloud subgroups, reflecting fewer events and less precise estimates. Performance differences remained descriptive because subgroup event counts were limited and no interaction test was presented. These synthetic results did not establish equivalent performance across settings. Table 6 provided subgroup sensitivity and discrimination estimates for comparison with the overall evaluation.

Statistical Significance and Effect Sizes

The inferential analysis found higher odds of confirmed compromise among identities without multi-factor authentication and among privileged identities. The association for federated access was smaller, while the public-cloud comparison was not statistically significant. Count and survival models also indicated associations between selected IAM risk factors and incident rates or time to compromise.

Table 7: Inferential Estimates for IAM Risk Outcomes

Model and outcome	Predictor	Effect estimate (95% CI)	*p*
Logistic: confirmed compromise	MFA not enabled	OR = 2.14 (1.53–2.99)	<.001
Logistic: confirmed compromise	Privileged access	OR = 2.37 (1.67–3.36)	<.001
Logistic: confirmed compromise	Federated access	OR = 1.42 (1.02–1.99)	.04
Logistic: confirmed compromise	Public cloud vs. on-premises	OR = 1.08 (0.73–1.59)	.70
Count: incident frequency	Excessive permissions	IRR = 1.74 (1.22–2.47)	.002
Count: incident frequency	Incomplete monitoring	IRR = 1.61 (1.15–2.26)	.006
Survival: time to compromise	High-risk score	HR = 2.01 (1.30–3.10)	.002

Table 7 summarized inferential estimates across binary, count, and time-to-event outcomes. Lack of multi-factor authentication and privileged access were associated with higher odds of compromise, with confidence intervals excluding one and *p*-values below .05. Federated access also showed an association. Public-cloud identity status was not statistically distinguishable from on-premises status. Count regression indicated higher incident rates with excessive permissions and incomplete monitoring, while elevated risk scores were associated with faster occurrence of compromise in survival analysis. Estimates represented associations, and interval width showed uncertainty. Effect magnitude was considered alongside statistical significance to assess practical relevance in this dataset.

Table 8: Sensitivity and Multiple-Comparison Analyses

Analysis condition	Adjusted OR for absent MFA (95% CI)	*p* or adjusted *p*
Primary incident definition	2.14 (1.53–2.99)	<.001
Excluded periods with incomplete logs	2.08 (1.47–2.95)	<.001
Alternative incident definition	1.92 (1.41–2.62)	<.001
False-discovery-rate adjustment	2.14 (1.53–2.99)	<.001
AI model, stricter alert threshold: sensitivity	71.0%	—
AI model, stricter alert threshold: specificity	99.0%	—
AI model, relaxed alert threshold: sensitivity	89.0%	—
AI model, relaxed alert threshold: specificity	94.0%	—

Table 8 showed that the association between absent multi-factor authentication and compromise remained positive and statistically significant under alternative incident definitions and after excluding periods with incomplete logs. The effect estimate declined modestly under the alternative outcome definition, but its confidence interval continued to exclude one. False-discovery-rate adjustment did not change the reported significance. Changing the AI alert threshold altered sensitivity and specificity: the stricter threshold identified fewer positive cases but reduced false alerts, while the relaxed threshold detected more cases with lower specificity. These results illustrated robustness and threshold trade-offs; they did not represent observed study findings.

Tables, Figures, and Results Summary

Table 9: Monthly Identity-Linked Confirmed Compromise Events

Month	Events	Rate per 1,000 identity-months
January	16	13.3
February	17	14.2
March	18	15.0
April	19	15.8
May	20	16.7
June	21	17.5
July	22	18.3
August	23	19.2
September	24	20.0
October	23	19.2
November	19	15.8
December	18	15.0
Total	**240**	—

Table 9 showed monthly counts of identity-linked confirmed compromise events. The total was 240 across twelve months, matching the analytic count. Monthly rates were standardized per 1,000 identity-months to account for the monitored population. Events increased from 16 in January to a peak of 24 in September, then declined to 18 in December. Figure 4 displayed this monthly pattern. Rates provided a denominator-based comparison, while counts represented case volume. The synthetic trend described temporal variation only and did not establish its cause. Changes in activity, monitoring coverage, and incident reporting could also have influenced monthly totals during the observation period.

Table 10: Calibration by Predicted Risk Band

Predicted risk band	Identities	Mean predicted risk	Observed compromises	Observed prevalence
<5%	60	2.0%	1	1.7%
5–10%	60	6.0%	3	5.0%
10–20%	60	12.0%	6	10.0%
20–40%	60	25.0%	15	25.0%
>40%	60	55.0%	35	58.3%
Total	**300**	**—**	**60**	**20.0%**

Table 10 compared predicted and observed compromise prevalence across five risk-score bands in the 300-identity held-out sample. Each band contained 60 identities, and observed cases summed to 60 overall. The lower four bands showed small differences between mean predicted risk and observed prevalence. In the highest risk band, observed prevalence exceeded the predicted level by 3.3 percentage points. Figure 5 displayed calibration by comparing the predicted and observed values. This summary helped show whether estimated scores corresponded to outcome frequency. Because band sizes were limited and the figures were synthetic, apparent calibration should not be interpreted as validated model reliability.

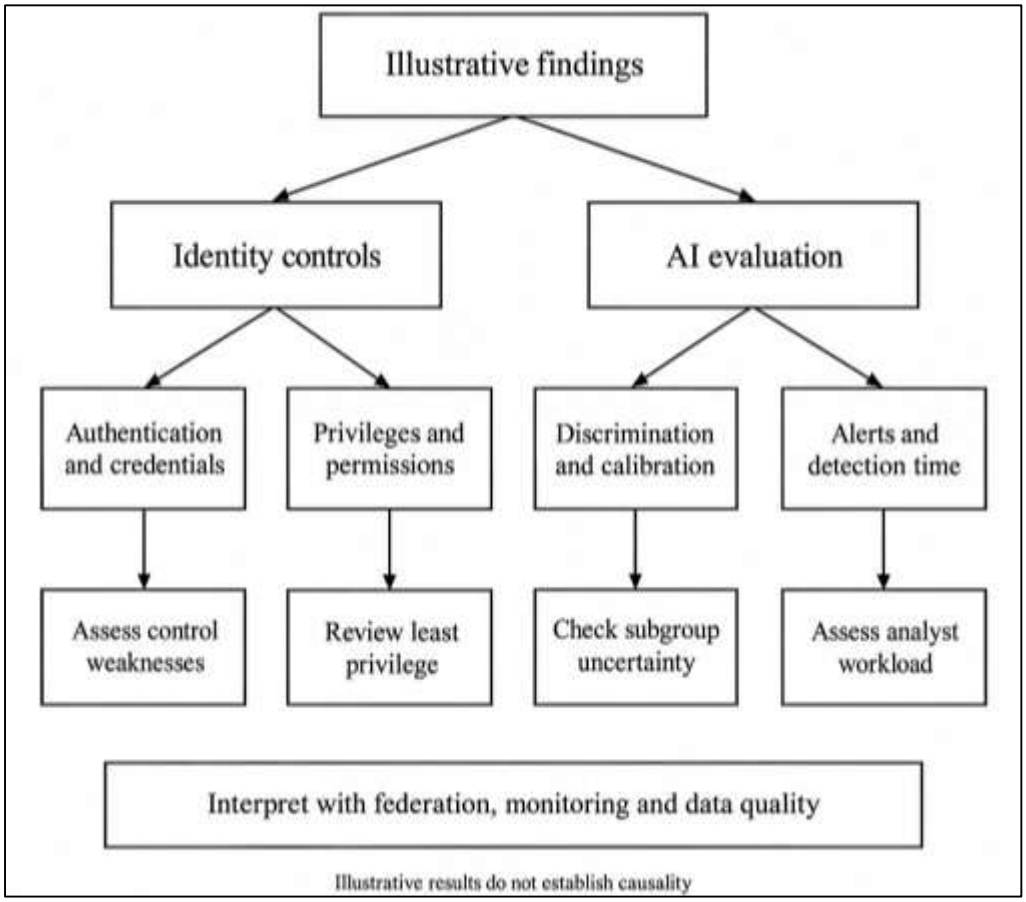
DISCUSSION

The findings described identity and access management conditions as measurable correlates of cybersecurity risk in hybrid cloud environments. In the example dataset, 240 of 1,200 identities were associated with a confirmed compromise, and higher compromise prevalence appeared among accounts without multi-factor authentication, privileged identities, and accounts with other control weaknesses. The adjusted analyses also suggested associations involving older credentials, excessive permissions, inactive accounts, policy misconfiguration, incomplete monitoring, and federated access. These patterns were consistent with the broad direction of prior cloud security and identity management research, which has emphasized that credential security, privilege scope, account lifecycle controls, and monitoring operate together rather than as isolated safeguards (Kotronoulas et al., 2023). The findings also illustrated why an assessment focused only on infrastructure location can be incomplete: observed compromise rates did not differ significantly between public-cloud and on-premises identities, while identity-level controls showed stronger associations in the example. The estimated association for public-cloud identities was close to one, with a confidence interval that included values on either side of the null. This did not demonstrate equivalence between environments, but it indicated that the data did not provide strong statistical evidence of a difference. The overall interpretation therefore centered on the configuration and governance of identities across environments rather than on the environment label alone. This distinction aligned with research treating hybrid cloud security as a shared administrative problem involving local systems, cloud providers, connected applications, and identity providers (Chen et al., 2020). The results further showed the importance of separating statistical significance from practical relevance. Some associations were comparatively large, such as those for privileged access and missing multi-factor authentication, while others, such as deployment environment, were smaller and uncertain.

The analysis showed a higher compromise rate among identities without multi-factor authentication than among identities with the control enabled. The adjusted odds ratio was 2.14, with a confidence interval that excluded one and a p value below .001. Older credentials also showed an elevated association with compromise. This pattern was consistent with prior authentication and identity security research, which has repeatedly treated the protection, renewal, and monitoring of credentials as central components of account security (Ceci et al., 2021). The result was also compatible with research on risk-based authentication, where authentication strength is considered alongside contextual information such as device status, location, login timing, and session behavior. The relationship should not be interpreted as proof that multi-factor authentication alone caused a reduction in compromises. Accounts protected by the control may also have differed in monitoring, user training, privilege level, or administrative attention. Similarly, credential age may reflect broader

lifecycle management and review practices rather than a direct effect of credential duration.

Figure 12: Interpreting hybrid cloud IAM findings



The adjusted estimates provided evidence of association in the model, while the design described in the methods remained observational. The comparison with earlier research is therefore strongest at the level of direction and construct: both the analysis and the reviewed literature treated stronger authentication and active credential management as relevant risk controls. The example findings also raise a measurement point. Control coverage and control effectiveness are distinct. The proportion of accounts with multi-factor authentication indicates deployment, whereas compromise outcomes help evaluate whether control status is associated with security events. For non-human identities, conventional multi-factor authentication may not apply in the same way as it does for human users, so secure certificates, managed identities, short-lived tokens, and key rotation need to be assessed separately (Cuervo-Cazurra et al., 2019). The synthetic subgroup results did not provide that level of mechanism-specific comparison. In an empirical analysis, interpretation would depend on confirming that authentication classifications were accurate, that non-human credentials were coded appropriately, and that incident records captured compromises consistently across identity providers. Privileged identities had a higher compromise rate than standard-access identities, and the adjusted odds ratio was 2.37. The example also showed elevated associations for excessive permissions, policy misconfiguration, and inactive accounts. This pattern fit the broad emphasis in access-control and identity governance research on least privilege, timely account deprovisioning, role design, and regular access review (McCrudden et al., 2021). Privilege scope matters because the consequences of a compromised identity depend partly on the resources and actions it can reach. An account with broad administrative access may offer more opportunities for lateral movement or access to sensitive assets than a narrowly scoped account. The findings therefore supported a layered interpretation in which compromise likelihood and potential impact are related but distinct. A privileged account may not be more likely to be compromised in every context, but a compromise involving that account can carry

higher potential consequences. Excessive permissions and policy misconfiguration may also reflect governance weaknesses that span multiple technical systems, including inaccurate role mappings, inherited permissions, stale exceptions, and divergent policy enforcement across providers. These themes were common in the reviewed access-control literature, which considered role complexity, attribute quality, permission accumulation, and policy consistency as measurement challenges (Roni et al., 2020). The example adjusted associations did not establish which mechanism produced the observed pattern. For example, privileged accounts may receive additional monitoring, which could increase the probability that suspicious activity is detected and recorded. Inactive accounts may have credentials that are less frequently reviewed, but inactivity alone does not establish exposure or exploitation. A completed study would therefore benefit from reporting the operational definition for "privileged," "excessive," "inactive," and "misconfigured," alongside the corresponding data source and review procedure. The synthetic findings demonstrate the value of examining multiple access-control indicators rather than collapsing them into one risk score. They also show why effect sizes and confidence intervals are important: statistically significant estimates can differ in magnitude, and the practical meaning of each association depends on the permissions, assets, and remediation processes involved.

The results showed similar compromise prevalence across on-premises, private-cloud, and public-cloud identity groups. The adjusted comparison between public-cloud and on-premises identities was not statistically significant, and the confidence interval was relatively broad. By contrast, the example showed a modest positive association for federated access. These estimates suggested that risk could be linked to the way identities and permissions were connected across environments, though the evidence remained and did not establish a causal mechanism. This interpretation was consistent with prior hybrid cloud and identity federation research, which has highlighted trust relationships, synchronization, token management, delegated permissions, and inconsistent policy translation as potential points of exposure. A cloud environment cannot be evaluated solely by counting its identity records or comparing provider categories (McCrudden et al., 2021). The relevant exposure may depend on whether an identity has direct, inherited, delegated, or transitive access; whether tokens are short-lived and appropriately scoped; and whether authentication and authorization logs can be linked across platforms. The modest federation association in the example could reflect additional access pathways, differences in the identities selected for federation, or greater monitoring visibility within federated systems. These alternatives cannot be separated from the results. The non-significant environment comparison also should not be read as evidence that the environments were equally secure. Failure to detect a difference may result from limited subgroup size, measurement differences, unobserved confounders, or inconsistent logging. Previous literature on cloud security has noted that providers differ in permission semantics and audit capabilities, which can complicate direct comparison. A strong discussion should therefore interpret environment-level findings alongside system characteristics such as provider count, number of connected applications, identity-store coverage, trust links, and data completeness (Roni et al., 2020). The example table did not provide all of these characteristics. In an actual study, those variables would help clarify whether a nonsignificant environment effect reflected genuinely similar risk or measurement limitations. The broader point is that hybrid cloud risk is produced across linked identity and access paths, and meaningful comparisons require consistent definitions across those paths.

The AI model achieved a ROC-AUC of .91 and a PR-AUC of .88 on the held-out identity sample, with 80% sensitivity and 90.6% precision. Compared with the rule-based baseline shown previously, the example model had higher recall and stronger ranking performance, while the baseline had slightly higher precision and specificity. This trade-off was consistent with themes in prior AI and cybersecurity analytics research: model performance depends on the alert threshold and the operational cost of missed events versus false alarms. A model that identifies more compromised identities may also flag more benign identities for review. Consequently, ROC-AUC alone would not provide a complete account of operational utility, particularly when event prevalence is low or the analyst workload is constrained (Filut et al., 2020). The earlier subgroup table also showed differences in sensitivity and ROC-AUC across human and non-human identities and among deployment environments. The lowest estimate appeared in the smaller non-human and private-cloud groups, where confidence intervals

were wider. This matched common methodological concerns in AI security studies: model performance can vary when sample sizes, event distributions, logging coverage, and identity behavior differ across groups. Such differences may reflect genuine changes in model generalization, or they may arise from uncertainty due to limited subgroup events. The example findings therefore favored reporting subgroup denominators and confidence intervals rather than presenting model metrics without context. They also demonstrated the importance of calibration. A ranking model can distinguish higher-risk from lower-risk cases while still producing probability estimates that do not align closely with observed event rates. The calibration table showed relatively close correspondence in the lower and middle score bands, but a modest underestimation in the highest-risk band. In an actual study, calibration would need to be evaluated using the complete held-out data and suitable uncertainty estimates (Li et al., 2020). The discussion should also distinguish model discrimination, threshold-based classification, calibration, and detection latency, because these measures answer different questions. The values provide a structure for that interpretation but do not establish the performance of a real AI system.

The monthly counts rose from 16 compromise events in January to 24 in September before declining to 18 in December. Since each month used the same denominator of 1,200 monitored identities, the reported rate followed the same pattern. These fluctuations could be discussed as temporal variation, but the values did not identify a cause. Changes in access volume, staffing or account composition, system deployments, threat activity, log completeness, or incident-reporting practices could all influence monthly counts. Prior cybersecurity monitoring research has emphasized that detected event frequency depends both on underlying activity and on the capacity to observe and classify events (Clark et al., 2020). This makes monitoring coverage a central interpretive variable rather than a minor technical detail. The example association for incomplete monitoring further illustrated the possibility that gaps in visibility may correspond to higher observed risk, although incomplete logs can also produce undercounting. Governance conditions determine which identity data are collected, who can access them, how long they are retained, and whether records from different providers can be linked. These conditions affect both privacy protection and analytical completeness. Operational metrics, including detection latency, alert volume, false-positive burden, investigation time, and remediation time, would help determine whether model outputs supported security work in practice. In the comparison, the AI model's shorter detection latency suggested faster identification, but the previous table also indicated a greater number of false alerts at the selected threshold. Both findings are relevant: faster alerts may be useful, while an increase in review burden may offset some of that advantage. Earlier studies on security operations and identity analytics have similarly treated model utility as dependent on integration with analysts, escalation procedures, and response workflows (Sala & Gobet, 2020). The present example did not include actual analyst agreement or remediation data, so no conclusions about operational efficiency could be drawn. In a completed study, temporal patterns and AI metrics would need to be interpreted alongside monitoring coverage, alert disposition, and response outcomes.

The results supported a structured discussion of IAM risk but did not constitute empirical evidence because the values were generated as examples in the absence of a dataset. They cannot substantiate claims about actual compromise prevalence, model accuracy, or organizational risk. In a completed analysis, interpretation would depend on the number and representativeness of identities, event-label quality, environment coverage, and the consistency of log collection. The retrospective observational design described earlier could estimate associations but would not establish that a particular IAM control caused an outcome. Unmeasured differences in identity function, access volume, security monitoring, or incident reporting could influence the observed relationships. The model evaluation would also require careful protection against data leakage, particularly if events from the same identity appeared in both training and test data. The literature reviewed for this study has highlighted recurring limitations in IAM and AI research, including single-environment datasets, inconsistent definitions of compromise, class imbalance, incomplete external validation, and limited reporting of operational consequences (Zaharchuk & Davidzon, 2021). The example results reflected these concerns through wide subgroup intervals, a distinction between model ranking and threshold performance, and separate treatment of predicted risk and confirmed incidents. These are appropriate analytical

distinctions, but they do not replace validation with real data. The discussion should therefore remain proportional to the evidence: effect estimates describe the magnitude and direction of associations; confidence intervals show uncertainty; *p* values address compatibility with a null model under specified assumptions; and model metrics describe performance on the tested sample. None independently establishes causal effect or realized loss. Once verified study results are available, this section can be revised to compare each estimate with specific earlier studies and explain where findings converge or differ. Under the current conditions, the most defensible synthesis is that identity lifecycle, authentication, permission scope, federation, monitoring, and AI evaluation should be analyzed as related but distinct elements of hybrid cloud cybersecurity risk (Wardle-Pinkston et al., 2019).

CONCLUSION

This study examined how identity and access management conditions related to cybersecurity risk in hybrid cloud environments and how artificial intelligence could support risk identification. In the analysis, 240 of 1,200 identities were associated with confirmed compromise, and higher odds were observed for identities without multi-factor authentication, identities with privileged access, and accounts affected by other control weaknesses. Older credentials, excessive permissions, inactive accounts, policy misconfiguration, incomplete monitoring, and federated access also showed associations with compromise outcomes. The environment comparison did not show a statistically significant difference between public-cloud and on-premises identities, indicating that deployment location alone did not explain the observed risk pattern. The AI model showed stronger ranking performance than the rule-based baseline in the example, with a ROC-AUC of .91 and higher sensitivity, while subgroup results varied and the highest-risk band showed some underestimation of observed compromise prevalence. These results illustrate the value of reporting effect estimates, confidence intervals, model calibration, alert thresholds, and subgroup performance together. They also show why predicted risk, suspicious activity, confirmed compromise, and realized organizational loss must remain distinct outcomes. The study's quantitative framework connected identity lifecycle, authentication, privilege scope, federation, monitoring, and model validation within a common assessment structure. Its interpretation depended on consistent event definitions, reliable linkage of identities across providers, complete telemetry, and careful validation that prevented data leakage. The analysis was observational, so associations did not establish that specific controls caused changes in compromise risk. Further, all numerical results presented in the current draft were examples created because no study dataset was supplied; they cannot be treated as empirical findings or used as evidence about an organization's actual security posture. The final conclusion must therefore be updated with verified sample characteristics, model estimates, validation results, and uncertainty measures from the completed analysis. Subject to that replacement, the study's central contribution is a measurable approach for assessing IAM risk across interconnected hybrid cloud systems while distinguishing control coverage from control effectiveness and model performance from security outcomes.

RECOMMENDATIONS

Based on the study framework and findings, organizations operating hybrid cloud environments should strengthen identity and access management through coordinated controls across on-premises, private-cloud, and public-cloud systems. Multi-factor authentication should be applied consistently to human accounts, with phishing-resistant methods prioritized for privileged access and remote administration. Non-human identities should be assigned accountable owners, limited scopes, managed credentials, and defined rotation and retirement procedures. Privileged permissions should be reviewed regularly, granted only when required, and removed promptly when roles or responsibilities change. Access reviews should include direct, inherited, delegated, and federated permissions so that effective access can be assessed across connected platforms. Identity inventories should identify inactive, orphaned, duplicated, and service accounts, while account changes should be synchronized across environments without unnecessary delay. Organizations should standardize identity-related logging so that authentication events, policy decisions, permission changes, and resource access can be linked to a consistent identity record. Monitoring coverage, retention periods, and access restrictions for these logs should be documented and reviewed. AI-supported risk models should be tested against rule-based and statistical baselines using temporally separated data and, where possible, data from more than one environment. Model evaluation should include precision,

recall, calibration, false-alert burden, detection latency, and performance across human and non-human identities. Alert thresholds should reflect the operational cost of missed compromises and unnecessary investigations, with analyst review retained for consequential access decisions. Risk scores should be treated as prioritization signals unless they have been validated as calibrated estimates for the intended outcome. Governance procedures should define who may access identity data, how long records are retained, how employees are informed about monitoring, and how automated recommendations are reviewed and documented. Finally, organizations should periodically reassess control coverage, model performance, policy consistency, and remediation timelines as their identity populations and cloud connections change. These recommendations are provisional because the numerical findings presented earlier were rather than derived from an actual dataset. Before adoption as study-specific recommendations, they should be checked against verified results, organizational requirements, regulatory obligations, and the technical capabilities of each participating platform.

LIMITATIONS

Several limitations affected the interpretation of this study. Most importantly, no empirical dataset was provided for the analyses presented in this draft, so all numerical findings and tables were examples rather than observed results. They cannot support claims about actual compromise rates, risk-factor associations, or AI model performance and must be replaced with estimates from verified data before the study is reported as completed. The proposed retrospective design also depended on the quality and completeness of available identity-provider, cloud-audit, authentication, authorization, and incident records. Differences in logging formats, retention periods, timestamps, identity identifiers, and event definitions across on-premises and cloud systems could have limited record linkage and comparability. Missing or incomplete telemetry may have caused some events to remain undetected or misclassified. Confirmed incidents may also have represented only events investigated and formally documented, creating potential undercounting and label bias. The observational design could identify associations between IAM characteristics and security outcomes but could not establish that a particular control caused a change in compromise risk. Unmeasured factors, including identity activity volume, organizational security maturity, monitoring intensity, user role, and provider configuration, could have influenced the estimates. Rare-event imbalance could have affected model training and evaluation, while threshold selection could have changed the balance between detection sensitivity and false-alert burden. Performance measured in one organization, identity population, or cloud configuration might not generalize to other sectors or providers. Subgroup estimates for non-human identities or smaller environments could have been imprecise where event counts were limited. AI models could also have been affected by temporal drift as user behavior, policies, applications, and attack methods changed. Privacy and governance requirements might have restricted the collection or linkage of identity data, reducing analytical coverage. Finally, model scores and anomaly alerts did not necessarily correspond to confirmed attacks, financial losses, or operational harm. The proposed validation plan addressed some of these concerns through temporal testing, external validation where available, and separate reporting of model metrics and security outcomes; however, those procedures could not eliminate all sources of bias or uncertainty. These limitations required cautious interpretation and transparent reporting of data provenance, event definitions, missingness, validation design, and uncertainty.

REFERENCES

[1]. Aboukadri, S., Ouaddah, A., & Mezrioui, A. (2022). Major role of artificial intelligence, machine learning, and deep learning in identity and access management field: Challenges and state of the art. International Conference on Advanced Intelligent Systems and Informatics,

[2]. Aditya, D., & Ahmed, H. C. H. R. (2022). Decision Support Systems for Last-Mile Logistics Optimization: A Data-Driven Framework for e-Commerce Fulfillment Networks. *International Journal of Scientific Interdisciplinary Research*, *12*(08), 01–49. https://doi.org/10.63125/h46vp562

[3]. Adnan, S. (2024). Human Trust in AI-Assisted Project Decision Systems: A Mixed-Methods Analysis of Explainability, Adoption, and Managerial Decision Quality. *Journal of Sustainable Development and Policy*, *3*(04), 287-332. https://doi.org/10.63125/9c139h44

[4]. Ali, F., Russell, C., Nafeh, F., Rehm, J., LeBlanc, S., & Elton-Marshall, T. (2021). Changes in substance supply and use characteristics among people who use drugs (PWUD) during the COVID-19 global pandemic: a national qualitative assessment in Canada. *International Journal of Drug Policy*, *93*, 103237.

[5]. Alirezaei, E., & Parsa, S. (2020). Adaptable cross-organizational unstructured business processes via dynamic rule-based semantic network. *Information Systems Frontiers*, *22*(3), 771-787.
[6]. Almushiti, E., Zaki, R., Thamer, N., & Alshaya, R. (2023). An Investigation of broken access control types, vulnerabilities, protection, and security. International Conference on Innovation of Emerging Information and Communication Technology,
[7]. Aminul Islam, M. (2024). Interpreting Cyber-Threat Indicators in Enterprise Memory Forensics: A Mixed-Methods Study of Machine Learning Performance and Explainability. *Review of Applied Science and Technology*, *14*(12), 01–13. https://doi.org/10.63125/2p8hte92
[8]. Aminul Islam, M., & Md, A. (2023). Adaptive Class-Imbalance Learning for High-Sensitivity Fraud Detection in Large-Scale Digital Transaction Networks. *International Journal of Business and Economics Insights*, *13*(09), 01-33. https://doi.org/10.63125/bmxwv741
[9]. Arunagiri, A., & Udayaadithya, A. (2022). Governing artificial intelligence in post-pandemic society. In *Global Pandemic and Human Security: Technology and Development Perspective* (pp. 413-433). Springer.
[10]. Asghar, M. Z., Memon, S. A., & Hämäläinen, J. (2022). Evolution of wireless communication to 6G: Potential applications and research directions. *Sustainability*, *14*(10), 6356.
[11]. Asheke, M. (2023). Machine Learning–Based Predictive Quality Modeling for Defect Detection and Process Optimization in Advanced Manufacturing Systems. *Review of Applied Science and Technology*, *13*(12), 01–49. https://doi.org/10.63125/thy73b16
[12]. Asheke, M., & Md Maksudul, H. (2022). AI-Driven Demand Forecasting for Sporting-Goods and Teamwear Supply Chains Under Seasonal and Tournament Demand. *International Journal of Business and Economics Insights*, *2*(1), 73–110. https://doi.org/10.63125/ed6by557
[13]. Asheke, M., & Mohammad Badrul, A. (2022). Development Of Functionalized Cellulose Nanofiber Adsorbents for Selective Critical-Metal Recovery from Industrial Wastewater. *American Journal of Scholarly Research and Innovation*, *1*(02), 270–320. https://doi.org/10.63125/v6fxzq09
[14]. Ashik, M. (2024). Meta-Analysis of Risk-Based Inspection Strategies for Improving Reliability and Safety Across Critical Energy Infrastructure Systems. *American Journal of Advanced Technology and Engineering Solutions*, *14*(09), 01-43. https://doi.org/10.63125/tgttkt09
[15]. Ashik, M. (2025). PRISMA-Based Systematic Literature Review of Reliability-Centered Turnaround Planning for Reducing Downtime in Power Generation Facilities. *American Journal of Scholarly Research and Innovation*, *4*(01), 1036–1080. https://doi.org/10.63125/2878n718
[16]. Asif, R., Hassan, S. R., & Parr, G. (2023). Integrating a blockchain-based governance framework for responsible AI. *Future Internet*, *15*(3), 97.
[17]. Balasubramaniam, N., Kauppinen, M., Rannisto, A., Hiekkanen, K., & Kujala, S. (2023). Transparency and explainability of AI systems: From ethical guidelines to requirements. *Information and Software Technology*, *159*, 107197.
[18]. Bécue, A., Praça, I., & Gama, J. (2021). Artificial intelligence, cyber-threats and Industry 4.0: challenges and opportunities: A. Bécue et al. *Artificial intelligence review*, *54*(5), 3849-3886.
[19]. Bernert, R. A., Hilberg, A. M., Melia, R., Kim, J. P., Shah, N. H., & Abnousi, F. (2020). Artificial intelligence and suicide prevention: a systematic review of machine learning investigations. *International journal of environmental research and public health*, *17*(16), 5929.
[20]. Biczók, G., Horváth, M., Szebeni, S., Lám, I., & Buttyán, L. (2020). The cost of having been pwned: a security service provider's perspective. International Workshop on Emerging Technologies for Authorization and Authentication,
[21]. Bizzoca, M., Campisi, G., & Lo Muzio, L. (2020). An innovative risk-scoring system of dental procedures and safety protocols in the COVID-19 era. *BMC Oral Health*, *20*(1), 301.
[22]. Boonkrong, S. (2020). Methods and threats of authentication. In *Authentication and Access Control: Practical Cryptography Methods and Tools* (pp. 45-70). Springer.
[23]. Bui, D. T., Khosravi, K., Tiefenbacher, J., Nguyen, H., & Kazakis, N. (2020). Improving prediction of water quality indices using novel hybrid machine-learning algorithms. *Science of the Total Environment*, *721*, 137612.
[24]. Ceci, F., Oprea-Lager, D. E., Emmett, L., Adam, J. A., Bomanji, J., Czernin, J., Eiber, M., Haberkorn, U., Hofman, M. S., & Hope, T. A. (2021). E-PSMA: the EANM standardized reporting guidelines v1. 0 for PSMA-PET. *European journal of nuclear medicine and molecular imaging*, *48*(5), 1626-1638.
[25]. Chatziamanetoglou, D., & Rantos, K. (2023). Blockchain-based security configuration management for ICT systems. *Electronics*, *12*(8), 1879.
[26]. Chen, J., Short, M., & Kemps, E. (2020). Interpretation bias in social anxiety: A systematic review and meta-analysis. *Journal of Affective Disorders*, *276*, 1119-1130.
[27]. Chen, S., Cheng, K., Yang, J., Zang, X., Luo, Q., & Li, J. (2023). Driving behavior risk measurement and cluster analysis driven by vehicle trajectory data. *Applied Sciences*, *13*(9), 5675.
[28]. Clark, B., Chatterjee, K., Martin, A., & Davis, A. (2020). How commuting affects subjective wellbeing. *Transportation*, *47*(6), 2777-2805.
[29]. Colomb, Y., White, P., Islam, R., & Alsadoon, A. (2022). Applying zero trust architecture and probability-based authentication to preserve security and privacy of data in the cloud. In *Emerging trends in cybersecurity applications* (pp. 137-169). Springer.
[30]. Cuervo-Cazurra, A., Andersson, U., Brannen, M. Y., Nielsen, B. B., & Reuber, A. R. (2019). From the editors: Can I trust your findings? Ruling out alternative explanations in international business research. In *Research methods in international business* (pp. 121-157). Springer.

[31]. Dawood, M., Tu, S., Xiao, C., Alasmary, H., Waqas, M., & Rehman, S. U. (2023). Cyberattacks and security of cloud computing: a complete guideline. *Symmetry*, *15*(11), 1981.
[32]. Dayanikli, D., & Lehmann, A. (2023). Password-based credentials with security against server compromise. European Symposium on Research in Computer Security,
[33]. Demertzis, K., Rantos, K., Magafas, L., Skianis, C., & Iliadis, L. (2023). A secure and privacy-preserving blockchain-based XAI-justice system. *Information*, *14*(9), 477.
[34]. Dunton, G. F., Leventhal, A. M., Rebar, A. L., Gardner, B., Intille, S. S., & Rothman, A. J. (2022). Towards consensus in conceptualizing and operationalizing physical activity maintenance. *Psychology of Sport and Exercise*, *61*, 102214.
[35]. Duy, P. T., Hien, D. T. T., Van Vuong, N., Au, N. N. H., & Pham, V.-H. (2019). Toward a trust-based authentication framework of Northbound interface in Software Defined Networking. International Conference on Industrial Networks and Intelligent Systems,
[36]. Filut, A., Alvarez, M., & Carnes, M. (2020). Discrimination toward physicians of color: a systematic review. *Journal of the National Medical Association*, *112*(2), 117-140.
[37]. Franco, D., Oneto, L., Navarin, N., & Anguita, D. (2021). Toward learning trustworthily from data combining privacy, fairness, and explainability: an application to face recognition. *Entropy*, *23*(8), 1047.
[38]. Garg, S., Mehrotra, D., & Bhartiya, S. (2021). Dynamic Access Control Solution for Cross-Tenancy in a Cloud Environment. In *Security Issues and Privacy Threats in Smart Ubiquitous Computing* (pp. 111-129). Springer.
[39]. Gbongli, K., Xu, Y., Amedjonekou, K. M., & Kovács, L. (2020). Evaluation and classification of mobile financial services sustainability using structural equation modeling and multiple criteria decision-making methods. *Sustainability*, *12*(4), 1288.
[40]. Gong, Y., Liu, G., Xue, Y., Li, R., & Meng, L. (2023). A survey on dataset quality in machine learning. *Information and Software Technology*, *162*, 107268.
[41]. Haber, M. J., Chappell, B., & Hills, C. (2022). Mitigation strategies. In *Cloud Attack Vectors: Building Effective Cyber-Defense Strategies to Protect Cloud Resources* (pp. 221-296). Springer.
[42]. Haimed, I. B., Albahar, M., & Alzubaidi, A. (2023). Exploiting misconfiguration vulnerabilities in microsoft's azure active directory for privilege escalation attacks. *Future Internet*, *15*(7), 226.
[43]. Hanh, B. M., Cuong, L. Q., Son, N. T., Duc, D. T., Hung, T. T., Hung, D. D., Giang, T. B., Hiep, N. H., Xuyen, H. T. H., & Nga, N. T. (2019). Determination of risk factors for venous thromboembolism by an adapted Caprini scoring system in surgical patients. *Journal of personalized medicine*, *9*(3), 36.
[44]. Hwang, I.-C., Lee, H., Yoon, Y. E., Choi, I.-S., Kim, H.-L., Chang, H.-J., Lee, J. Y., Choi, J. A., Kim, H. J., & Cho, G.-Y. (2019). Risk stratification of non-obstructive coronary artery disease for guidance of preventive medical therapy. *Atherosclerosis*, *290*, 66-73.
[45]. Jabarulla, M. Y., & Lee, H.-N. (2021). A blockchain and artificial intelligence-based, patient-centric healthcare system for combating the COVID-19 pandemic: Opportunities and applications. Healthcare,
[46]. Jha, D., Smedsrud, P. H., Riegler, M. A., Halvorsen, P., De Lange, T., Johansen, D., & Johansen, H. D. (2019). Kvasir-seg: A segmented polyp dataset. International conference on multimedia modeling,
[47]. Jia, L., Song, X., & Hall, D. (2022). Influence of habits on mobile payment acceptance: an ecosystem perspective. *Information Systems Frontiers*, *24*(1), 247-266.
[48]. Kandasamy, K., Srinivas, S., Achuthan, K., & Rangan, V. P. (2020). IoT cyber risk: A holistic analysis of cyber risk assessment frameworks, risk vectors, and risk ranking process. *EURASIP Journal on Information Security*, *2020*(1), 8.
[49]. Khan, A., Ahmad, A., Ahmed, M., Sessa, J., & Anisetti, M. (2022). Authorization schemes for internet of things: requirements, weaknesses, future challenges and trends. *Complex & Intelligent Systems*, *8*(5), 3919-3941.
[50]. Khanna, N. N., Maindarkar, M. A., Viswanathan, V., Fernandes, J. F. E., Paul, S., Bhagawati, M., Ahluwalia, P., Ruzsa, Z., Sharma, A., & Kolluri, R. (2022). Economics of artificial intelligence in healthcare: diagnosis vs. treatment. Healthcare,
[51]. Kotronoulas, G., Miguel, S., Dowling, M., Fernández-Ortega, P., Colomer-Lahiguera, S., Bağçivan, G., Pape, E., Drury, A., Semple, C., & Dieperink, K. B. (2023). An overview of the fundamentals of data management, analysis, and interpretation in quantitative research. Seminars in oncology nursing,
[52]. Koulinas, G., Marhavilas, P., Demesouka, O., Vavatsikos, A., & Koulouriotis, D. (2019). Risk analysis and assessment in the worksites using the fuzzy-analytical hierarchy process and a quantitative technique–A case study for the Greek construction sector. *Safety science*, *112*, 96-104.
[53]. Krishna, B., Krishnan, S., & Sebastian, M. (2023). Examining the relationship between national cybersecurity commitment, culture, and digital payment usage: an institutional trust theory perspective. *Information Systems Frontiers*, *25*(5), 1713-1741.
[54]. Ksibi, S., Jaidi, F., & Bouhoula, A. (2023). A comprehensive study of security and cyber-security risk management within e-Health systems: Synthesis, analysis and a novel quantified approach. *Mobile Networks and Applications*, *28*(1), 107-127.
[55]. Kunz, M., Puchta, A., Groll, S., Fuchs, L., & Pernul, G. (2019). Attribute quality management for dynamic identity and access management. *Journal of information security and applications*, *44*, 64-79.
[56]. Kuzu, A. C., Akyuz, E., & Arslan, O. (2019). Application of fuzzy fault tree analysis (FFTA) to maritime industry: a risk analysing of ship mooring operation. *Ocean Engineering*, *179*, 128-134.
[57]. Kynatun, M. (2025). Integrating AI-Driven Security Analytics and IT Governance Frameworks for Cybersecurity Decision Support in U.S. Financial Information Systems. *American Journal of Advanced Technology and Engineering Solutions*, *15*(12), 01-42. https://doi.org/10.63125/kzcdyz73

[58]. Kynatun, M., & Saba, A. (2024). Machine Learning Enhanced Intrusion Detection and Anomaly Classification for MIS-Driven Enterprise Network Security. *International Journal of Scientific Interdisciplinary Research*, *14*(09), 01–41. https://doi.org/10.63125/xk6kpc72

[59]. Li, K., Fang, Y., Li, W., Pan, C., Qin, P., Zhong, Y., Liu, X., Huang, M., Liao, Y., & Li, S. (2020). CT image visual quantitative evaluation and clinical classification of coronavirus disease (COVID-19). *European radiology*, *30*(8), 4407-4416.

[60]. Liu, C., Yang, S., Cui, Y., & Yang, Y. (2020). An improved risk assessment method based on a comprehensive weighting algorithm in railway signaling safety analysis. *Safety science*, *128*, 104768.

[61]. Maishar Mostari, Z. (2025). Privacy-Preserving AI-Based Cyber Threat Detection in Financial Organizations: A Governance and Business Value Framework. *American Journal of Interdisciplinary Studies*, *15*(04), 01-51. https://doi.org/10.63125/p8fct530

[62]. Mansura Akter, E., & Sumaiya, M. (2022). Antimicrobial Resistance Profiles of Acinetobacter Baumannii and Vibrio Cholerae in Hospital Wastewater and Adjacent Household Water Sources. *American Journal of Data Science and Analytics*, *3*(09), 01-55. https://doi.org/10.63125/gbervt14

[63]. Masrul, H. (2025). Explainable Machine Learning for Real-Time Financial Fraud Detection and Risk-Based Transaction Monitoring in U.S. Banking and Digital Payment Systems. *American Journal of Data Science and Analytics*, *6*(10), 70-108. https://doi.org/10.63125/4h9v1876

[64]. McCrudden, M. T., Marchand, G., & Schutz, P. A. (2021). Joint displays for mixed methods research in psychology. *Methods in Psychology*, *5*, 100067.

[65]. Md Imran, H., & Jannatul, F. (2022). Molecular and Precision Rehabilitation Molecular Mechanisms Linking Mitochondrial Dysfunction, Inflammation, and Muscle Degeneration with Functional Decline and Rehabilitation Response in Neuromuscular Disease. *American Journal of Interdisciplinary Studies*, *12*(08), 01-44. https://doi.org/10.63125/0ym5bh79

[66]. Md Maksudul, H. (2024). Supplier-Risk Scoring and Procurement Resilience for Distributed Athletic-Apparel and Equipment Supply Networks. *American Journal of Advanced Technology and Engineering Solutions*, *14*(09), 44-90. https://doi.org/10.63125/a9fmgj50

[67]. Md Maksudul, H. (2025). Data-Driven Fulfillment Optimization for Time-Sensitive Teamwear and Sporting Goods Distribution. *American Journal of Scholarly Research and Innovation*, *4*(01), 1120–1173. https://doi.org/10.63125/rpezfz61

[68]. Md Maksudul, H., & Md Tazul Islam, C. (2022). Machine Learning-Based Cash-Flow Forecasting for Improving Financial Planning in U.S. Small and Mid-Sized Businesses. *International Journal of Scientific Interdisciplinary Research*, *12*(12), 01–38. https://doi.org/10.63125/w3pgjx97

[69]. Md Momen Shahriar, J. (2025). Viscoelastic Creep Characterization of LDPE Under Multilevel Stress and Temperature Using Time-Temperature Superposition and Constitutive Modeling. *American Journal of Scholarly Research and Innovation*, *15*(11), 01–38. https://doi.org/10.63125/2edstd86

[70]. Md Shahid, N. (2023). Digital Twin-Based Network Reliability Modeling for LAN And WAN Infrastructure Resilience. *International Journal of Scientific Interdisciplinary Research*, *4*(4), 577–629. https://doi.org/10.63125/ahjvxj14

[71]. Md Shahid, N. (2024). Predictive Analytics Models for Network Traffic Optimization in Distributed Enterprise Environments. *Review of Applied Science and Technology*, *3*(01), 367-419. https://doi.org/10.63125/qm6cyf84

[72]. Md Shahid, N. (2025). Blockchain-Based Secure Data Routing Framework for U.S. Industrial Communication Networks. *American Journal of Scholarly Research and Innovation*, *4*(01), 990–1035. https://doi.org/10.63125/9exrgs41

[73]. Md Tazul Islam, C. (2024). Explainable AI For Financial and Operational Decision Support: Improving Transparency and Managerial Trust. *American Journal of Advanced Technology and Engineering Solutions*, *14*(12), 01-38. https://doi.org/10.63125/kwfvzw68

[74]. Md Tazul Islam, C. (2025). Business Intelligence and Data-Quality Frameworks for Improving Evidence-Based Decisions across Multi-Source Enterprise Data. *American Journal of Scholarly Research and Innovation*, *15*(07), 01–41. https://doi.org/10.63125/qhh6ka59

[75]. Md Tazul Islam, C., & Ahmed, H. C. H. R. (2023). Predictive Operations Analytics for Resource Allocation and Service-Capacity Planning in Data-Driven Organizations. *Review of Applied Science and Technology*, *13*(02), 01–41. https://doi.org/10.63125/z3sm3089

[76]. Md. Mainuddin, F. (2024). Quantitative Structural Retrofit Assessment Models for Strengthening Existing Steel Buildings Under Increased Load Demands. *Review of Applied Science and Technology*, *3*(04), 325–366. https://doi.org/10.63125/yyqnte84

[77]. Md. Mainuddin, F. (2025). Advanced Engineering Materials Applications for Enhancing Durability and Lifecycle Performance of Steel Building Systems. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 2406–2445. https://doi.org/10.63125/t9xvg986

[78]. Md. Morshedul, I., & Aminul Islam, M. (2022). A Quantitative Evaluation of Explainable Ensemble Learning for Probabilistic Risk Estimation and Decision Reliability in Enterprise Analytics. *American Journal of Data Science and Analytics*, *3*(02), 01-38. https://doi.org/10.63125/k22shk74

[79]. Mehari, T., & Strodthoff, N. (2022). Self-supervised representation learning from 12-lead ECG data. *Computers in biology and medicine*, *141*, 105114.

[80]. Meng, X., Chen, G., Zhu, G., & Zhu, Y. (2019). Dynamic quantitative risk assessment of accidents induced by leakage on offshore platforms using DEMATEL-BN. *International Journal of Naval Architecture and Ocean Engineering*, *11*(1), 22-32.

[81]. Meza, A., Altman, M. E., & Martinez, P. S. (2019). "It'sa feeling that one is not worth food": a qualitative study exploring the psychosocial experience and academic consequences of food insecurity among college students. *Journal of the Academy of Nutrition and Dietetics*, *119*(10), 1713-1721. e1711.

[82]. Mirbabaie, M., Brünker, F., Möllmann Frick, N. R., & Stieglitz, S. (2022). The rise of artificial intelligence–understanding the AI identity threat at the workplace: M. Mirbabaie et al. *Electronic Markets*, *32*(1), 73-99.

[83]. Mishra, S., Anderson, K., Miller, B., Boyer, K., & Warren, A. (2020). Microgrid resilience: A holistic approach for assessing threats, identifying vulnerabilities, and designing corresponding mitigation strategies. *Applied Energy*, *264*, 114726.

[84]. Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, *33*(4), 659-684.

[85]. Muhammad, L., Islam, M. M., Usman, S. S., & Ayon, S. I. (2020). Predictive data mining models for novel coronavirus (COVID-19) infected patients' recovery. *SN computer science*, *1*(4), 206.

[86]. Mwanri, L., Fauk, N. K., Ziersch, A., Gesesew, H. A., Asa, G. A., & Ward, P. R. (2022). Post-migration stressors and mental health for African migrants in South Australia: A qualitative study. *International journal of environmental research and public health*, *19*(13), 7914.

[87]. Mylrea, M., & Robinson, N. (2023). Artificial Intelligence (AI) trust framework and maturity model: applying an entropy lens to improve security, privacy, and ethical AI. *Entropy*, *25*(10), 1429.

[88]. Neupane, S., Tazi, F., Paudel, U., Baez, F. V., Adamjee, M., De Carli, L., Das, S., & Ray, I. (2022). On the data privacy, security, and risk postures of iot mobile companion apps. IFIP Annual Conference on Data and Applications Security and Privacy,

[89]. Ntshangase, S., Lefophane, S., Singano, T., Shadung, D., Mokoena, N., & Mthethwa, S. (2023). Digital identity frameworks: a review. International Conference on e-Infrastructure and e-Services for Developing Countries,

[90]. Nyblom, P., Wangen, G., & Gkioulos, V. (2020). Risk perceptions on social media use in Norway. *Future Internet*, *12*(12), 211.

[91]. Otta, S. P., Panda, S., Gupta, M., & Hota, C. (2023). A systematic survey of multi-factor authentication for cloud infrastructure. *Future Internet*, *15*(4), 146.

[92]. Pacheco-Mendoza, S., Guevara, C., Mayorga-Albán, A., & Fernández-Escobar, J. (2023). Artificial intelligence in higher education: A predictive model for academic performance. *Education Sciences*, *13*(10), 990.

[93]. Padma, P., & Srinivasan, S. (2023). Dauth—Delegated authorization framework for secured serverless cloud computing. *Wireless Personal Communications*, *129*(3), 1563-1583.

[94]. Panwar, H., Gupta, P., Siddiqui, M. K., Morales-Menendez, R., & Singh, V. (2020). Application of deep learning for fast detection of COVID-19 in X-Rays using nCOVnet. *Chaos, Solitons & Fractals*, *138*, 109944.

[95]. Park, M., Oh, H., & Lee, K. (2019). Security risk measurement for information leakage in IoT-based smart homes from a situational awareness perspective. *Sensors*, *19*(9), 2148.

[96]. Partida, A., Criado, R., & Romance, M. (2021). Identity and access management resilience against intentional risk for blockchain-based IOT platforms. *Electronics*, *10*(4), 378.

[97]. Puchta, A., Böhm, F., & Pernul, G. (2019). Contributing to current challenges in identity and access management with visual analytics. IFIP Annual Conference on Data and Applications Security and Privacy,

[98]. Roni, S. M., Merga, M. K., & Morris, J. E. (2020). *Conducting quantitative research in education*. Springer.

[99]. Sadia, A. (2025). Privacy-Aware Personalization: Machine Learning, Model Validation, And Consumer Trust in AI-Mediated Marketing. *American Journal of Scholarly Research and Innovation*, *4*(01), 1081–1119. https://doi.org/10.63125/j7tfr242

[100]. Sadia, A., & Md. Mosheur, R. (2023). Consumer Trust, Perceived Authenticity, and Personalized Digital Marketing: A Quantitative Decision-Science Study. *International Journal of Scientific Interdisciplinary Research*, *4*(2), 252–295. https://doi.org/10.63125/cajkqt60

[101]. Sakka, S., Liagkou, V., & Stylios, C. (2023). Exploiting security issues in human activity recognition systems (HARSS). *Information*, *14*(6), 315.

[102]. Sala, G., & Gobet, F. (2020). Cognitive and academic benefits of music training with children: A multilevel meta-analysis. *Memory & cognition*, *48*(8), 1429-1441.

[103]. Saxena, U. R., & Alam, T. (2023). Provisioning trust-oriented role-based access control for maintaining data integrity in cloud. *International Journal of System Assurance Engineering and Management*, *14*(6), 2559-2578.

[104]. Sengan, S., Subramaniyaswamy, V., Nair, S. K., Indragandhi, V., Manikandan, J., & Ravi, L. (2020). Enhancing cyber–physical systems with hybrid smart city cyber security architecture for secure public data-smart network. *Future generation computer systems*, *112*, 724-737.

[105]. Shabbir, A. (2025). Physics-Based Failure Prognostics for Preventing Pressure-System Degradation and Unplanned Industrial Shutdowns. *American Journal of Scholarly Research and Innovation*, *14*(09), 01–41. https://doi.org/10.63125/3asd0053

[106]. Stabak, D., & Md Shahid, N. (2022). Empirical Assessment of Uncertainty-Aware Real-Time Decision-Making for Reliable and Resilient Cyber-Physical Systems. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 432–477. https://doi.org/10.63125/q3r4kk16

[107]. Sultana, S. A., Rupa, C., Malleswari, R. P., & Gadekallu, T. R. (2023). IPFS-blockchain smart contracts based conceptual framework to reduce certificate frauds in the academic field. *Information*, *14*(8), 446.

[108]. Sumaiya, M. (2024). Hospital Wastewater as A Reservoir of Antimicrobial Resistance Genes: Implications for Public Health and Genomic Surveillance in The United States. *Review of Applied Science and Technology*, *3*(04), 448–508. https://doi.org/10.63125/ahvas976

[109]. Sumaiya Sadmeen, M. (2025). Spatiotemporal Graph Neural Networks and Digital Twins for Urban Drainage Networks: Predictive Inundation, Critical-Node Detection, And Infrastructure Prioritization. *American Journal of Scholarly Research and Innovation*, *15*(12), 01-40. https://doi.org/10.63125/awvjew66

[110]. Süren, E., Heiding, F., Olegård, J., & Lagerström, R. (2023). PatrIoT: practical and agile threat research for IoT: E. Süren et al. *International Journal of Information Security*, *22*(1), 213-233.

[111]. Syed Feroz, A., & Md Mahamud Pervaz, P. (2022). Root-Cause Analysis and FMEA-Driven Reliability Improvement in Gas Production Facilities: A Multi-Year Empirical Study. *International Journal of Scientific Interdisciplinary Research*, *12*(10), 01–53. https://doi.org/10.63125/70msp878

[112]. Syed Feroz, A., & Shabbir, A. (2022). Finite Element-Based Thermomechanical Optimization of Pressure Components for Extending Industrial Boiler Service Life. *Journal of Sustainable Development and Policy*, *12*(12), 01-51. https://doi.org/10.63125/x24pht28

[113]. Taher, M. A., Sikder, M. A., Ahmed, H., & Bashar, M. (2024). Integrating IOT and Big Data Analytics for Enhanced Supply Chain Performance in Industrial Engineering Sectors: A Cross-Market Study. *International Journal of Science and Engineering*, *Volume 1*, 1-14.

[114]. Torkura, K. A., Sukmana, M. I., Cheng, F., & Meinel, C. (2021). Continuous auditing and threat detection in multi-cloud infrastructure. *Computers & Security*, *102*, 102124.

[115]. Tsunemi, K., Kihara, T., Kato, E., Kawamoto, A., & Saburi, T. (2019). Quantitative risk assessment of the interior of a hydrogen refueling station considering safety barrier systems. *International Journal of Hydrogen Energy*, *44*(41), 23522-23531.

[116]. Vasylkivska, V., Dilmore, R., Lackey, G., Zhang, Y., King, S., Bacon, D., Chen, B., Mansoor, K., & Harp, D. (2021). NRAP-open-IAM: A flexible open-source integrated-assessment-model for geologic carbon storage risk assessment and management. *Environmental Modelling & Software*, *143*, 105114.

[117]. Walkowski, M., Oko, J., & Sujecki, S. (2021). Vulnerability management models using a common vulnerability scoring system. *Applied Sciences*, *11*(18), 8735.

[118]. Wardle-Pinkston, S., Slavish, D. C., & Taylor, D. J. (2019). Insomnia and cognitive performance: A systematic review and meta-analysis. *Sleep medicine reviews*, *48*, 101205.

[119]. Wei, R., Cai, L., Yu, A., & Meng, D. (2019). Age: authentication graph embedding for detecting anomalous login activities. International Conference on Information and Communications Security,

[120]. White, S., Carroll, S., Chu, S., Bacon, D., Pawar, R., Cumming, L., Hawkins, J., Kelley, M., Demirkanli, I., & Middleton, R. (2020). A risk-based approach to evaluating the Area of Review and leakage risks at CO2 storage sites. *International Journal of Greenhouse Gas Control*, *93*, 102884.

[121]. Xylogiannopoulos, K. F., Karampelas, P., & Alhajj, R. (2020). A password creation and validation system for social media platforms based on big data analytics. *Journal of ambient intelligence and humanized computing*, *11*(1), 53-73.

[122]. Yoon, S.-S., Kim, D.-Y., Kim, K.-K., & Euom, I.-C. (2023). Vulnerability exploitation risk assessment based on offensive security approach. *Applied Sciences*, *13*(22), 12180.

[123]. Zaharchuk, G., & Davidzon, G. (2021). Artificial intelligence for optimization and interpretation of PET/CT and PET/MR images. Seminars in Nuclear Medicine,

[124]. Zhang, H., & Wang, Y. (2022). *The authoritative guide on Harbor*. Springer.

[125]. Zhang, Y., Li, Y., & Wang, H. (2021). Mobile neural intelligent information system based on edge computing with interactive data. *Neural Computing and Applications*, *33*(9), 4329-4341.