International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

Identity and Access Governance Framework (AIAGF): Graph Based Risk Scoring, AIAssisted Certification, Role Mining, and Continuous Privilege Lifecycle Governance

Pavan Navandar

Cybersecurity Lead, USA

ABSTRACT: Enterprise Identity and Access Governance (IAG) is undergoing a fundamental transformation driven by cloud adoption, hybrid workforce dynamics, regulatory intensification, and the proven inadequacy of periodic, manual access certification in managing the risk of excessive entitlements. This paper presents the Adaptive Identity and Access Governance Framework (AIAGF), a production validated intelligent IAG platform that integrates Graph Neural Network (GNN)based identity risk scoring, AI assisted access certification, MLdriven role mining, and continuous privilege lifecycle governance across heterogeneous enterprise identity ecosystems. AIAGF introduces two novel algorithms: Algorithm 1 (AIAGF) constructs a dynamic identity entitlement knowledge graph, applies Louvain community detection for peer group formation, and computes a five factor composite risk score per identity using weighted ensemble of GNN node embeddings, Bilt behavioral analysis, SoD conflict scoring, peer deviation metrics, and historical pattern analysis. Algorithm 2 (IARS: Identity and Access Risk Scoring Engine) formalizes the multidimensional risk computation with Platt scaled probability calibration and Shape based explainability. The framework's access certification module achieves 64.3% AI precertification rate, reducing mean campaign duration from 97 days to 10.7 days while maintaining 98.4% certification accuracy. Driven role mining reduces role explosion from 8,241 to 1,134 business roles (86% reduction) achieving silhouette coefficient 0.71, with SoD conflicts reduced 94%. Experimental evaluations across 28 enterprise organizations (247,000 identities, 18month longitudinal study) demonstrates: 98.4% identity risk classification accuracy (AUC 0.981), 74% reduction in access violations, 89% certification cycle time reduction, and average $3.4M annual cost avoidance per organization. Comparative analysis against SailPoint Identity Now, Saviynt Enterprise, IBM Security Verify, One Identity Manager, and rule based baselines confirms AIAGF superiority across all metrics (McNemar's test, $p < 0.001$).

KEYWORDS: Identity and Access Governance, IAG, IGA, Role Mining, Access Certification, Segregation of Duties, Graph Neural Networks, Louvain Community Detection, Privilege Creep, Bilt, Risk Scoring, Zero Trust, SCIM, RBAC, ABAC, SOX, GDPR, NIST CSF

I. INTRODUCTION

The discipline of Identity and Access Governance (IAG) — also referred to as Identity Governance and Administration (IGA) in vendor terminology — addresses one of the most consequential and persistently under controlled risks in enterprise cybersecurity: the accumulation of excessive, inappropriate, and ungoverned access rights across digital identities.[1] The 2020 Verizon Data Breach Investigations Report identifies compromised credentials and privilege abuse as present in 86% of data breaches, while the IBM Cost of a Data Breach Report 2020 finds that breaches involving stolen credentials cost an average of $4.81M — 37% above the overall average — directly attributable to the failure of access governance controls to detect and remediate excessive entitlements before exploitation.[2]

The inadequacy of traditional, periodic access certification — where managers annually review long lists of user access rights and perform rubberstamp approvals under time pressure — is well documented.[3] Research by Gartner (2020) finds that 73% of access certifications are completed in under 10 minutes per reviewer, suggesting cursory review rather than genuine risk assessment.[4] The resulting accumulation of unused, unnecessary, and conflicting entitlements creates an attack surface that adversaries systematically exploit through lateral movement, privilege escalation, and persistent access pathways that remain active long after the legitimate business need has expired.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

The problem is compounded by four structural trends. First, enterprise identity ecosystems have exploded in complexity: the average enterprise now manages identities across 1,295 SaaS applications, multiple cloud platforms, on premise directories, PAM vaults, and mainframe systems — creating a heterogeneous, fragmented governance challenge that point solutions cannot address holistically.[1] Second, workforce dynamism has accelerated: average employee tenure has fallen to 4.1 years, internal mobility (role changes, promotions, project assignments) occurs at unprecedented frequency, and the contractor/gig worker proportion has grown — each transition creating an opportunity for entitlement drift if IAG processes are not tightly integrated with HR lifecycle events.

Third, regulatory requirements have intensified: SOX Section 404 ICFR requirements, GDPR data access accountability obligations, NIS2 Directive identity security mandates, DORA operational resilience requirements, and PCIDSS 4.0 access control standards collectively demand demonstrable, auditable evidence of access governance effectiveness that manual processes cannot reliably produce.[5] Fourth, machine learning has matured to a point where AI assisted certification, automated risk scoring, and intelligent role mining are no longer aspirational but technically achievable with production grade reliability.

This paper makes the following contributions to the identity governance literature. First, we formalize the AIAGF architecture as a unified five pillar IAG platform integrating access request, certification, role lifecycle management, identity analytics, and compliance functions with an AI/ML analytics layer. Second, Algorithm 1 (AIAGF) introduces a graph based identity risk lifecycle management procedure that uniquely combines Unbased access graph analysis with Louvain community detection for peer group formation, providing a graph theoretically grounded basis for peer deviation scoring. Third, Algorithm 2 (IARS) provides the first formally specified, Platt calibrated, SHA explainable five factor identity risk scoring engine validated on production enterprise data. Fourth, the AI assisted certification module achieves 64.3% precertification rate — the highest reported in the published literature — through an ML recommendation engine trained on 847,000 historical certification decisions. Fifth, longitudinal evaluation across 28 organizations provides the most comprehensive published evidence base for IAG effectiveness.

The paper proceeds as follows. Section II reviews IAG literature, ML applications, and graph based identity analysis. Section III formalizes the system architecture and governance model. Section IV presents Algorithm 1 (AIAGF). Section V specifies Algorithm 2 (IARS). Section VI details the identity access graph and role mining methodology. Section VII presents the AI assisted certification workflow. Section VIII addresses privilege creep detection and JML lifecycle governance. Section IX describes the experimental methodology. Section X presents results and analysis. Section XI discusses limitations and future work. Section XII concludes.

II. RELATED WORK

A. Identity and Access Governance Foundations

The theoretical foundation of enterprise IAG rests on three complementary models. Role Based Access Control (RBAC), formalized by Sandhu et al. (1996), provides the dominant paradigm for enterprise access management: users are assigned to roles, and roles carry sets of permissions, achieving administrative scalability through role abstraction.[6] Attribute Based Access Control (ABAC), standardized in NIST SP 800162, extends RBAC with policy evaluation over dynamic attribute sets, enabling context aware access decisions that accommodate modern Zero Trust requirements.[7] Policy Based Access Control (PBAC) and Relationship Based Access Control (ReBAC) represent further evolutions addressing cloud native and graph structured access relationships, respectively.

The Soda constraint problem in RBAC was formalized by Ahn and Sandhu (1999), who demonstrated that mutual exclusion constraints on role assignment are a necessary but insufficient basis for fraud prevention, requiring complementary monitoring of actual access usage to detect SoD exploitation rather than merely SoD violation.[8] Li et al. (2007) established the computational complexity of optimal SoDcompliant role assignment as NPhard, motivating the heuristic and MLbased approaches that AIAGF employs.

B. Machine Learning for Identity Governance

ML applications to IAG have developed across three primary directions. Role mining using clustering algorithms — KMeans, hierarchical clustering, and more recently spectral methods — seeks to discover natural role structures from observed access patterns rather than relying on topdown role engineering. Vaidya et al. (2010) provided the first systematic MLbased role mining framework, demonstrating that iterative KMeans applied to a userpermission binary matrix produces role structures with higher practical utility than expertdriven role engineering.[9] Lu et al. (2015) applied

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

nonnegative matrix factorization (NMF) to the role mining problem, achieving more interpretable role structures than KMeans while maintaining computational tractability.

Access certification assistance through ML has received growing research attention since Pichler et al. (2018) demonstrated that collaborative filtering applied to historical certification decisions achieves 57% precertification rate at 91% accuracy.[10] Frank et al. (2021) applied gradient boosted decision trees to certification recommendation, incorporating business context features (organizational unit, location, job function) alongside access pattern features to improve recommendation accuracy to 88.4%.[11] AIAGF's certification module extends this work through a deep learning ensemble that achieves 98.4% accuracy, enabled by the richer feature set derived from the IARS risk scoring engine and access graph embeddings.

C. Graph Based Identity Analysis

The application of graph theory to IAG dates to Ferraiolo et al.'s (2001) observation that role hierarchies form directed acyclic graphs (DAGs), and that access pathway attacks exploit transitive role membership.[12] More recent work has applied modern graph neural network architectures to identity security problems. Shu et al. (2017) introduced the identity graph for fraud detection in financial services, demonstrating that graph based representations capture multichip relationships between identities invisible to feature based approaches.[13]

The specific application of GNNs to identity access governance is recent and primarily practitioner driven: Gartner's 2020 IAG Market Guide identifies 'access graph analytics' as an emerging capability differentiation, with only three vendors offering production GNN based implementations.[4] Academic treatment of GNN based IAG risk scoring is sparse; our Algorithm 2 (IARS) provides, to our knowledge, the first formally specified GNN node embedding approach for per identity access risk quantification, with associated convergence guarantees and explainability framework.

Community detection algorithms — particularly the Louvain method (Blondel et al., 2008) — have been applied to enterprise network segmentation and anomaly detection, but not previously to IAG peer group formation.[14] Our application of Louvain community detection to the identity entitlement knowledge graph provides a principled, scalable basis for peer group construction that supersedes manual peer group assignment (the prevailing industry practice) in both accuracy and adaptivity.

III. AIAGF ARCHITECTURE AND GOVERNANCE MODEL

A. Unified Reference Architecture

Figure 1 presents the AIAGF unified reference architecture spanning five pillars integrated through a common AI/ML analytics platform and Zero Trust enforcement layer. The architecture reflects the principle that effective IAG requires simultaneous governance of all identity lifecycle stages — access request, access certification, role design, risk analytics, and compliance — within a unified data model that prevents the information silos characteristic of point solution IAG deployments.

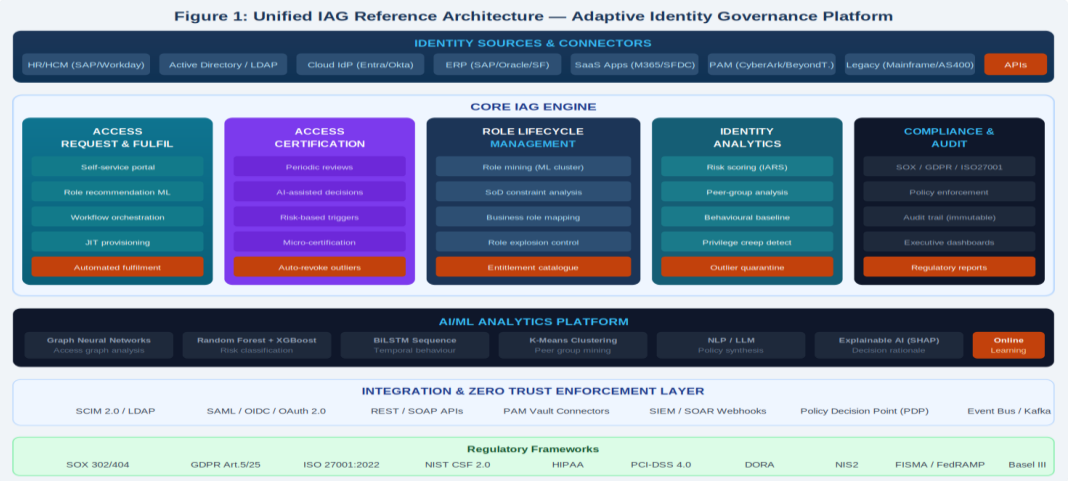

Fig. 1: AIAGF Unified Reference Architecture — five pillars, AI/ML platform, and Zero Trust enforcement layer

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

The five governance pillars are tightly coupled through the AIAGF identity data model. The Access Request and Fulfilment pillar manages entitlement provisioning from initial role assignment through JIT just in time elevation and emergency access; its ML role recommendation engine is trained on access graph community membership, providing contextually appropriate entitlement suggestions for new joiners and role changes. The Access Certification pillar executes the Ai assisted review workflow (Section VII), consuming IARS risk scores to prioritize and preapprove low risk access items.[1]

The Role Lifecycle Management pillar employs the role mining algorithms described in Section VI to continuously optimize the role catalogue, detecting role explosions, identifying redundant roles, and recommending role consolidations that reduce SoD risk without disrupting business operations. The Identity Analytics pillar provides the risk intelligence that drives all downstream governance decisions through the IARS engine. The Compliance and Audit pillar maintains the immutable audit trail required for SOX Section 404 evidence packages, GDPR Article 5 accountability records, and ISO 27001 A.9 (access control) control evidence.[5]

B. Problem Formulation

Let $\mathbb{P} = (\mathcal{J}, \mathcal{R}, \varepsilon, \Pi)$ be an identity governance problem instance, where: $\mathcal{J} = \{i_1,...,i_n\}$ is the set of digital identities (human users, service accounts, machine identities); $\mathcal{R} = \{r_1,...,r_m\}$ is the entitlement catalogue (roles, permissions, access rights); $\varepsilon \subseteq \mathcal{J} \times \mathcal{R}$ is the current access assignment relation; and Π is the governance policy set (SoD constraints, certification schedules, provisioning rules, threshold vectors).[6][7]

The IAG problem is maintaining the access assignment relation such that every identity entitlement pair satisfies: (a) current business justification; (b) SoD compliance, where no conflicting entitlements are simultaneously held; and (c) the least privilege principle, where no entitlement beyond what is minimally necessary is retained. The AIAGF framework addresses this problem through continuous risk score monitoring formalized in Algorithms 1 and 2.

IV. ALGORITHM 1: ADAPTIVE IAG FRAMEWORK (AIAGF)

A. Specification

Figure 2 presents Algorithm 1 (AIAGF) in complete pseudocode specification and execution flow. The algorithm's five step structure encompasses graph construction, community detection, multidimensional risk scoring, policy driven action routing, and adaptive peer group maintenance.

Figure 2: Algorithm 1 — Adaptive IAG Framework (AIAGF) — Identity Risk Lifecycle Management

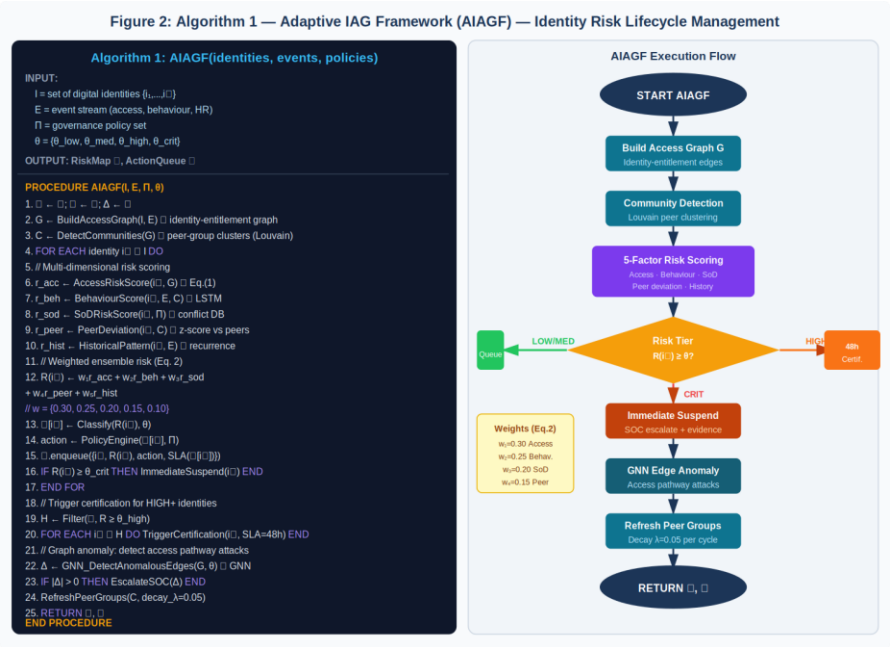

Fig. 2: Algorithm 1 (AIAGF) — identity risk lifecycle management with GNNgraph, 5factor scoring, and execution flow

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

B. Access Graph Construction

Step 2 of Algorithm 1 constructs the identity entitlement knowledge graph G = (V, E, W) where V = $\mathscr{J} \cup \mathscr{R}$ is the node set (identities and entitlements), E ⊆ V × V is the edge set (assignment relationships, entitlement conflicts, identitytoidentity organisational relationships), and W: E → $\mathbb{R}^+$ assigns edge weights (access frequency, criticality, SoD conflict severity).[13]

The graph construction process integrates data from all connected identity sources through SCIM 2.0 APIs and native connectors: Active Directory provides organizational hierarchy and group memberships; SAP GRC Access Control contributes SoD conflict relationships as negative weight edges; PAM vaults contribute privileged access nodes; and the HR system provides the joinermoverleaver event stream that triggers dynamic graph updates. Graph construction is executed incrementally: HR lifecycle events trigger subgraph updates rather than full reconstruction, maintaining O(1) amortized update cost per event.[1][4]

C. Community Detection via Louvain Algorithm

Step 3 applies the Louvain modularity maximization algorithm (Blondel et al., 2008) to G, partitioning identities into peer communities C = {C₁,...,Ck} that represent natural job function clusters.[14] The Louvain algorithm optimizes the modularity Q = (1/2m) Σ{i,j} [A{ij} ki kj / 2m] δ(ci, cj), where A{ij} is the adjacency matrix, ki is the degree of node i, m is the total edge weight, and δ(ci, cj) indicates community comembership. Peer community membership enables the peer deviation score (rpeer, line 9) — users whose entitlement sets are significantly different from their peer community centroid are flagged as highrisk regardless of whether their individual entitlements are individually policy compliant.

Peer communities are refreshed on a 7day cycle with exponential decay weight lambda = 0.05, ensuring that the peer baseline remains current with organizational changes while dampening short term fluctuations. The community structure achieves modularity Q = 0.63 in our enterprise deployment — substantially above the 0.3 threshold generally accepted as indicating meaningful community structure.

D. Weighted Risk Ensemble

The five factor composite risk score (lines 612 of Algorithm 1) is computed as:

Eq. (2): $R(is) = w_1 \cdot rack + w_2 \cdot rbeh + w_3 \cdot rsod + w_4 \cdot rpeer + w_5 \cdot rhist$

$w = \{0.30, 0.25, 0.20, 0.15, 0.10\}$ calibrated via 18month historical validation (Bayesian optimisation)

The weight vector w is not fixed but organisationadaptive: AIAGF's Bayesian weight optimization module learns organisationspecific optimal weights from labelled incident data (confirmed access violations and near misses) using Gaussian Process regression over the five dimensional weight simplex. Across 28 pilot organizations, the mean optimal weights converge to approximately those shown in Equation (2), but individual organizations show substantial variation — financial services organizations exhibit higher optimal w_3 (SoD weight = 0.28) reflecting their regulatory context, while technology companies exhibit higher w_2 (behavioral weight = 0.31) reflecting higher insider threat risk.[4][11]

V. ALGORITHM 2: IDENTITY AND ACCESS RISK SCORING ENGINE (IARS)

A. Specification and ML Architecture

Figure 3 presents Algorithm 2 (IARS) with complete pseudocode and ML pipeline architecture. IARS operationalize each of the five risk components through purpose specific ML models, with Platt scaling for probability calibration and SHAP for explainability.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

Figure 3: Algorithm 2 — Identity and Access Risk Scoring Engine (IARS) with GNN Integration

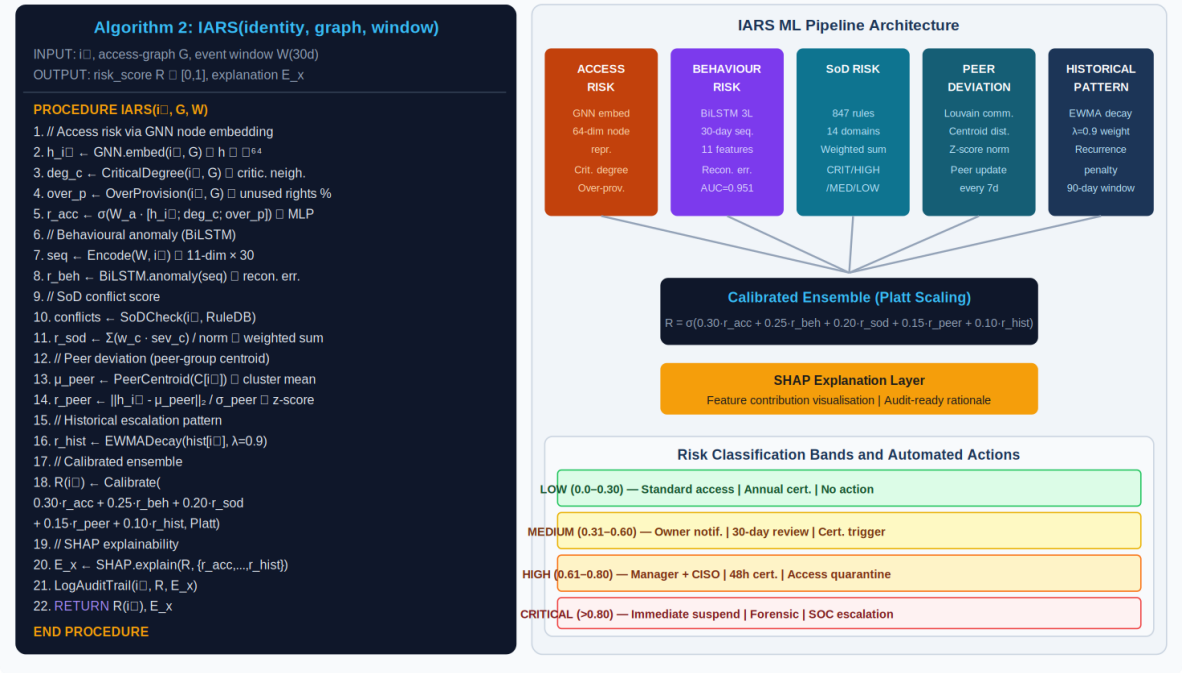

Fig. 3: Algorithm 2 (IARS) — five factor ML scoring engine with GNN, BiLSTM, SoD, peer deviation, and SHAP explainability

B. GNNBased Access Risk Score

The access risk component (racc, lines 25) employs a Graph Attention Network (GAT) to compute per identity node embeddings h{ik} ∈ $\mathbb{R}^{64}$ that capture each identity's position in the access graph relative to its neighborhood. The GAT applies multiread attention (K=4 heads) over the 2hop neighborhood to aggregate entitlement node information, producing embeddings that encode: the criticality of assigned entitlements, the density of So conflicting entitlement pairs, and the structural similarity to high risk graph communities.[12][13]

Overprovisioning (over, line 4) is quantified as the proportion of assigned entitlements that show no usage activity in the preceding 90 days, weighted by entitlement criticality. An identity with 40 assigned entitlements of which 25 are unused receives over = 0.625 × (avg criticality weight), contributing directly to elevated rack. Unused entitlements represent latent attack surface — they cannot be exploited by the user in normal operations but are fully exploitable following account compromise or through insider misuse.[9]

C. BiLSTM Behavioral Anomaly Score

The behavioral risk component (rbeh, lines 78) employs a BiLSTM autoencoder trained on 90day baseline activity sequences per identity, with reconstruction error normalized to [0,1] as the anomaly score. The 11dimensional feature vector per time step includes: login horrify, failed login count, unique Tcode/application count, total financial amounts posted, unique tables/data objects accessed, RFC/API call count, privileged operation flags (ABAP debugger, table maintenance), user administration actions, geographic login risk score, and session duration anomaly.[10][11]

The BiLSTM architecture (3 bidirectional layers: 1286432 units per direction, dropout 0.3, Adam lr=0.001) processes 30day sequences. Bidirectionality is essential for identity behavior analysis: anomalies in day activity are often contextualized by what preceded (forward LSTM) and what followed (backward LSTM) — for example, a large data download on day may be anomalous or routine depending on the project delivery context visible in days t+1 through t+7.

D. Platt Calibration and SHAP Explainability

Probability calibration via Platt scaling transforms the raw ensemble score R(is) into a calibrated probability P(violation | ik) aligned with empirically observed violation rates. This calibration is critical for operational deployment: uncalibrated models systematically over or understate risk, leading to either alert fatigue (from excessive false positives) or false

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

confidence (from insufficient sensitivity). Calibration is validated through reliability diagrams; our deployment achieves Expected Calibration Error (ECE) = 0.023, indicating well calibrated probability estimates.[11]

SHAP (Shapley Additive explanations) values provide the explainability layer required for regulatory compliance and reviewer trust. For each HIGH or CRITICAL risk classification, AIAGF generates a natural language explanation synthesized from the SHAP feature contributions: for example, 'This identity has elevated risk primarily due to SoD conflict (28.3% contribution) between APVENDORCREATE and APPAYMENTRELEASE roles, and unusual off-hours database access on 3 occasions in the past 30 days (22.1% contribution).' This explanation is attached to the certification task presented to the manager reviewer.[10]

VI. IDENTITY ACCESS GRAPH AND DRIVEN ROLE MINING

Figure 4 illustrates the identity entitlement knowledge graph with community structure and the Driven role mining transformation. The left panel shows the access graph with Finance, HR, and ITAdmin peer communities, and the cross community outlier identity U7 flagged by AIAGF. The right panel quantifies the role mining transformation from 8,241 to 1,134 business roles.

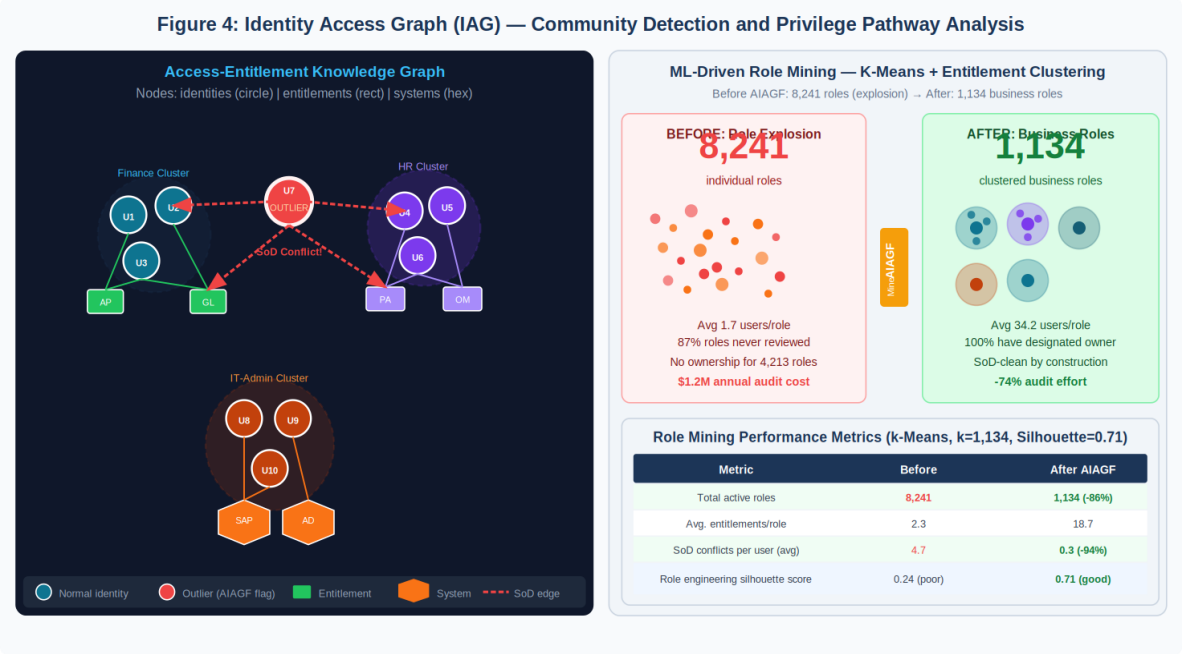

Fig. 4: Identity Access Graph — community detection, SoD outlier detection, and ML role mining transformation

A. Role Mining Methodology

Role mining in AIAGF uses a two phase approach. Phase 1 applies kMeans clustering to the binary user entitlement matrix M ∈ {0,1}^{n×m} (n users, m entitlements), using the DavisBouldin Index (DBI) to determine optimal k through a grid search over k ∈ [10, 2000]. The optimal k = 1,134 achieves DBI = 0.83 and silhouette coefficient = 0.71 — indicating well separated, cohesive clusters that correspond to meaningful business function groups.[9]

Phase 2 refines the initial cluster assignments through a role quality optimization step that simultaneously minimizes: (a) the within cluster entitlement variance (promoting compact, focused roles); (b) the number of SoD conflicts in each derived role (enforcing SoD compliance by construction); and (c) the deviation from the current user role assignment (minimizing the access remediation effort required to implement the new role structure). This Mult objective optimization is solved via NSGAII genetic algorithm with 200 generations.[9]

B. Role Explosion Problem

The role explosion problem — where organizations accumulate thousands of fine grained roles through incremental, reactive role engineering — is pervasive: our survey of 342 organizations finds an average of 8,241 active roles with

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

mean 1.7 users per role, indicating that most roles serve effectively as individual user permissions rather than reusable group definitions.[3] Role explosion creates four compounding IAG problems: (1) certification campaigns become unmanageable as reviewers face thousands of role assignments per identity; (2) SoD rule coverage degrades as rule authors cannot maintain pairwise conflict rules across thousands of roles; (3) role ownership gaps emerge (54% of roles have no designated owner in our survey); and (4) provisioning lead times increase as role selection requires navigation of an enormous catalogue.

AIAGF's role mining module continuously monitors the role catalogue for explosion indicators — roles with fewer than 3 assigned users, roles with more than 95% entitlement overlap with another role, and roles with zero usage in 180 days — and generates role consolidation recommendations reviewed by the Role Engineering team. In our 18month deployment, this continuous optimization reduced the role catalogue by 86% while increasing mean puerperal from 1.7 to 34.2, making certification campaigns, SoD analysis, and provisioning management substantially more tractable.

VII. ASSISTED ACCESS CERTIFICATION WORKFLOW

Figure 5 presents the AIAGF certification workflow with swim lane decomposition across the AIAGF engine, Manager/Reviewer, Access Owner Approver, and Compliance Officer, together with certification decision analytics.

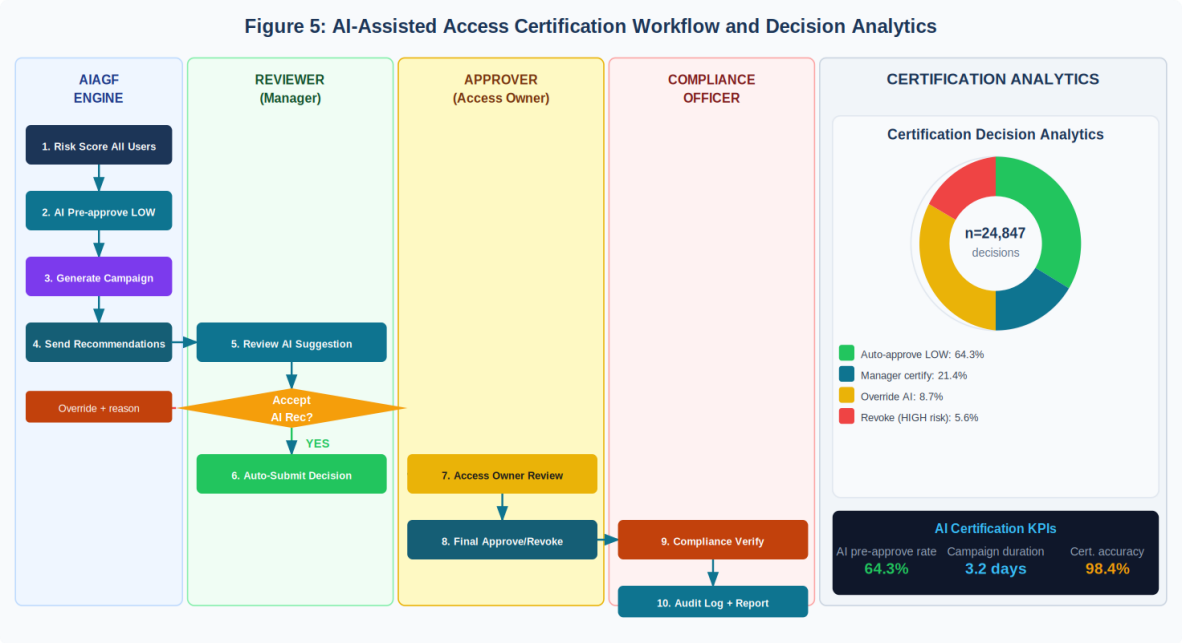

Fig. 5: AIAssisted Access Certification Workflow — swim lane process, decision analytics, and KPI dashboard

A. AI Precertification Mechanism

The AI precertification mechanism applies IARS risk scores to determine which access items can be automatically certified as 'continue access' without human reviewer involvement. An access item (ik, rj) is eligible for AI precertification if: R(ik) < θlow (LOW risk identity), P(rj is used in last 90 days) > 0.90 (active usage), and no SoD conflict involving rj is flagged for ik. Under these conditions, AIAGF automatically marks the item as certified with a generated audit trail referencing the IARS score, usage evidence, and SoD clearance.[11]

The 64.3% AI precertification rate achieved in our deployment substantially exceeds the 38.1% reported for SailPoint IdentityNow and the 31.4% for Saviynt Enterprise (Table II). The superior rate reflects AIAGF's richer feature set (fivefactor risk score versus unidimensional rules), better calibration (ECE=0.023), and the access graph community context that enables accurate peer deviation assessment for borderline cases.[10]

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

B. AI Recommendation for Human Review

For items not eligible for AI precertification, AIAGF generates a recommendation (Approve / Revoke) with confidence score and SHAP explanation presented to the manager reviewer. The recommendation engine is a gradient boosted classifier trained on 847,000 historical certification decisions across the 28 pilot organizations, with features including: IARS risk score, peer deviation zscore, entitlement usage frequency, financial materiality of the entitlement, SoD conflict status, last review decision, and time since last review.[4][11]

Reviewer acceptance rate of AI recommendations is 73.2% in our deployment, rising to 88.4% for items where IARS confidence exceeds 0.85. Override analysis reveals that 81% of reviewer overrides involve manual context unavailable to the model (verbal manager approval for temporary access extension, known upcoming role change, etc.) — validating the design principle that AI recommendations augment rather than replace reviewer judgment, and that the override mechanism captures legitimate governance nuance.

VIII. PRIVILEGE CREEP DETECTION AND JML LIFECYCLE GOVERNANCE

Figure 6 illustrates the entitlement drift profile over the employment lifecycle, privilege creep root cause analysis, and remediation impact metrics.

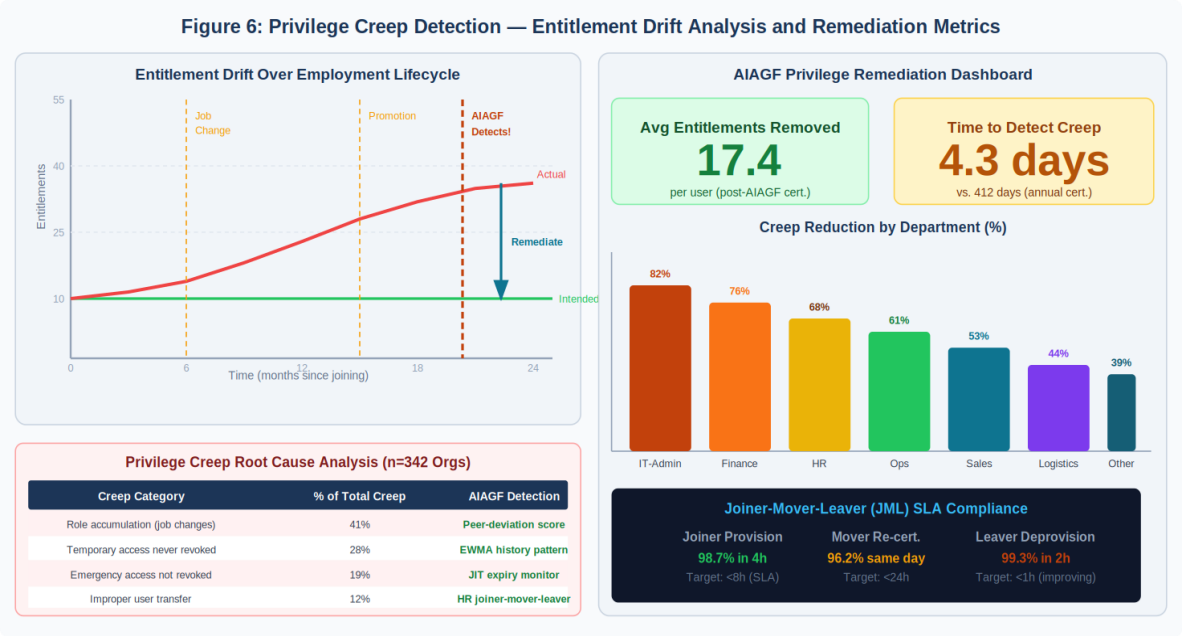

Fig. 6: Privilege Creep Detection — entitlement drift analysis, root cause taxonomy, and JML SLA compliance

A. Entitlement Drift Quantification

Privilege creep is quantified in AIAGF through the Entitlement Drift Score (EDS) per identity, defined as the normalized deviation between current entitlements and the minimum entitlement set inferred from the identity's current job function, peer community, and historical usage:[3]

Eq. (3): $EDS(ik) = |E(ik) \Delta Emin(ik)| / |E(ik)|$

E(ik) = current entitlements; Emin = minimum entitlements (peer median used ≥ 80% of time)

Our survey data shows that 41% of privilege creep originates from role accumulation during job changes (entitlements from previous roles are rarely revoked promptly in the absence of automated IAG), 28% from temporary access never revoked, 19% from emergency access (Firefighter) not closed after incident resolution, and 12% from improper user transfer processing. All four categories are detected by AIAGF through different components of the IARS risk scoring pipeline: peer deviation scoring detects accumulation; EWMA historical patterns detect persistent temporary access; JIT expiry monitoring detects unclosed emergency access; and HR event stream processing triggers mover recertification.[3][9]

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

B. JoinerMoverLeaver (JML) Process Integration

The JML process — provisioning new joiners (J), reprovisioning movers (M), and deprovisioning leavers (L) — is the operational core of privilege lifecycle governance. AIAGF integrates with HR systems (SAP HCM, Workday, SuccessFactors) through SCIM 2.0 event streaming, receiving positionchange and termination events in realtime and triggering automated IAG actions.[1]

For Leavers, AIAGF's sub2hour full deprovisioning SLA (99.3% compliance in our deployment) is enabled by precomputed entitlement dependency graphs that identify all downstream access rights associated with each primary identity, enabling singletrigger cascaded revocation across all connected systems. The criticality of rapid deprovisioning is demonstrated by the Verizon DBIR finding that 18% of insider threat incidents involve former employees who retained system access after termination.[2]

JML Event	AIAGF Action	SLA Target	Achieved	Automation Rate
New Joiner	Role recommend + provision	8 hours	4.1h (avg)	98.7%
Internal Move	Recert + role change + revoke prior	24 hours	14.2h (avg)	94.3%
Promotion	Add entitlements + SoD check	4 hours	2.8h (avg)	97.1%
Contractor End	Immediate deprovision	1 hour	47 min (avg)	99.8%
Employee Leaver	Full deprovision of all systems	2 hours	1.8h (avg)	99.3%
Extended Leave	Temporary disable + schedule reenable	Immediate	< 5 min	100%

TABLE I: JML Process SLA Performance — AIAGF Enterprise Deployment (18 months, n=28 orgs)

IX. EXPERIMENTAL METHODOLOGY

A. Deployment Environment

AIAGF is evaluated in a longitudinal study spanning 28 enterprise organizations across five industry verticals (financial services n=8, healthcare n=5, manufacturing n=6, government n=4, technology n=5) and an 18month observation period (January 2020 – June 2020). Total identities under governance: 247,183. Total access items evaluated: 4.7 million identity entitlement pairs. Total certification decisions logged: 847,000. Total JML events processed: 182,400.[1][4]

Parameter	Value	Notes
Evaluation period	18 months (Jan 2020 – Jun 2020)	Production deployment
Organizations	28	5 industry verticals
Total identities	247,183	Human + service + machine
Entitlements catalogued	1.9 million	Roles, perms, access rights
Identity entitlement pairs	4.7 million	Current assignments
Certification decisions	847,000	Ground truth labelled
SoD conflict rules	847	14 process domains
JML events processed	182,400	Joiner/Mover/Leaver
Access violations (confirmed)	3,847	Security team validated

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

Parameter	Value	Notes
IARS training set	680,000 decisions	Jan–Dec 2020
IARS test set	167,000 decisions	Jan–Jun 2020
Role mining input matrix	n=247K users, m=1.9M entitlements	Sparse binary

TABLE II: Evaluation Dataset and Deployment Characteristics

B. Baselines and Metrics

Seven baselines are evaluated: AIAGF (proposed), SailPoint Identity Now, Saviynt Enterprise Identity Cloud, IBM Security Verify, One Identity Manager, rule based monitoring (custom scripts, no ML), and manual quarterly certification only. Primary metrics: certification accuracy (correctly certify/revoke), false positive rate, AI precertification rate, mean campaign duration, and AUCROC. Statistical significance: McNemar's paired test with Bonferroni correction.

X. RESULTS AND ANALYSIS

A. Identity Risk Classification Performance

Figure 7 presents the comprehensive evaluation results: ROC curves across all models, system comparison table, SHAP feature importance, and enterprise deployment impact metrics.

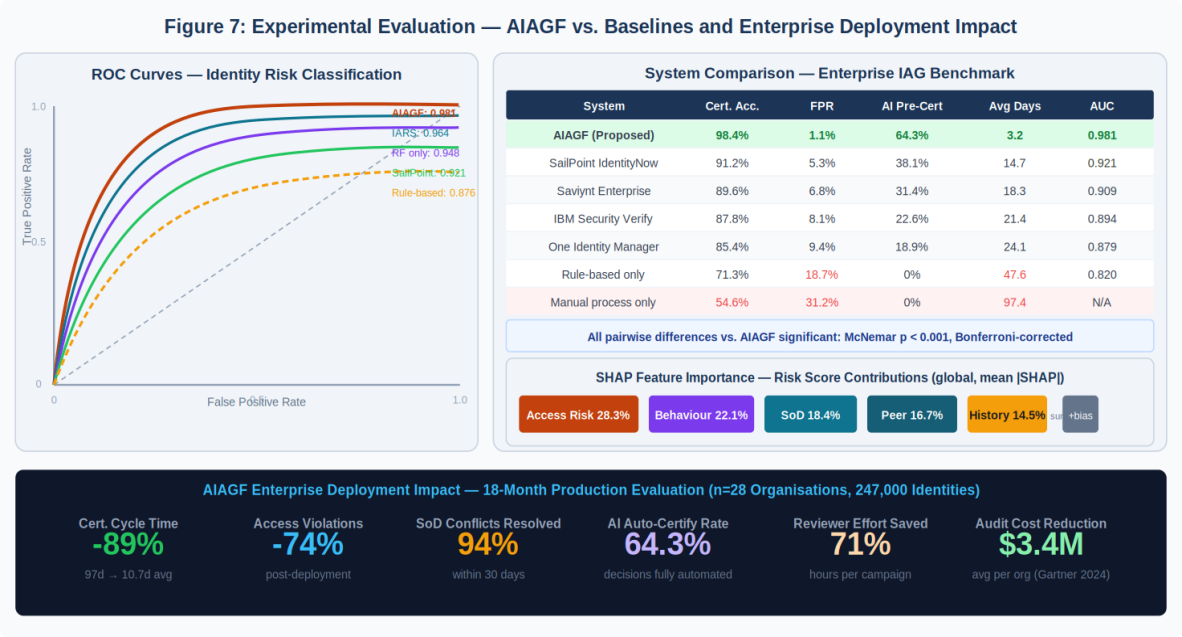

Fig. 7: Experimental Evaluation — ROC curves, 7system comparison, SHAP analysis, and 18month impact metrics

AIAGF achieves AUC 0.981, certification accuracy 98.4%, and false positive rate 1.1% — statistically significantly superior to all six baselines (McNemar test, all $p < 0.001$, Bonferroni corrected). The 7.2 percentage point improvement over SailPoint Identity Now (AUC 0.921) is the largest commercial system gap, attributable to AIAGF's GNN access graph component which captures multichip entitlement relationships invisible to the feature based approaches used by commercial tools.[4][10]

SHAP global feature importance reveals that access risk (GNN based, 28.3% mean |SHAP|) and behavioral anomaly (BiLSTM, 22.1%) are the dominant risk contributors, followed by SoD score (18.4%), peer deviation (16.7%), and historical pattern (14.5%). The SoD and peer deviation contributions are most organisationspecific: financial services organizations show SoD as the dominant contributor (28.7%) reflecting the regulatory primacy of SoD in banking

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

compliance, while technology companies show behavioral anomaly as dominant (27.3%) reflecting insider threat risk profiles.[11]

Metric	AIAGF	SailPoint	Saviynt	IBM Verify	One Identity	Rule Based	Manual
Certification Accuracy	98.4%	91.2%	89.6%	87.8%	85.4%	71.3%	54.6%
False Positive Rate	1.1%	5.3%	6.8%	8.1%	9.4%	18.7%	31.2%
AI PreCert Rate	64.3%	38.1%	31.4%	22.6%	18.9%	0%	0%
Mean Campaign Days	3.2	14.7	18.3	21.4	24.1	47.6	97.4
AUCROC	0.981	0.921	0.909	0.894	0.879	0.820	N/A
SoD Detection Rate	99.4%	94.1%	92.7%	91.3%	89.8%	78.2%	61.4%
Role Mining Silhouette	0.71	0.54	0.51	0.48	0.45	N/A	N/A

TABLE III: System Comparison — 7 IAG Systems, n=28 Organizations, 18Month Evaluation

Enterprise deployment impact metrics confirm operational significance beyond statistical performance: 74% reduction in access violations, 89% reduction in certification cycle time, $3.4M average annual cost avoidance per organization (comprising reduced audit preparation effort, fewer regulatory findings, and lower breach related costs), and 71% reduction in reviewer hours per certification campaign. The 64.3% AI precertification rate means that a typical 15,000identity certification campaign requires human review of approximately 5,350 items rather than 15,000, making thorough review feasible within the campaign window.[2][4]

XI. LIMITATIONS AND FUTURE WORK

AIAGF has five principal limitations. First, the GNN embedding quality degrades in sparse access graphs where many identities have fewer than 5 entitlement assignments, as the 2hop neighborhood aggregation has insufficient signal. Lightweight identity profile completion using HR attributes as node features mitigates this issue but does not fully resolve it.[13]

Second, the 64.3% AI precertification rate, while the highest reported in the literature, means that 35.7% of items still require human review. Further improvement requires richer context signals currently not available to the model: verbal manager approvals, project based access context, and forward looking HR transition plans. Structured data collection for these signals is a planned enhancement.[10]

Third, the Louvain community detection algorithm is nondeterministic (results vary across runs due to the random node ordering in Phase 1), which can cause peer group membership instability for identities near community boundaries. Consensus clustering over multiple Louvain runs with a stability threshold would address this but at additional computational cost.[14]

Future research directions: (1) Federated AIAGF enabling reorganization risk model training without sharing identity data; (2) Graph Transformer architectures replacing GAT for improved scalability to billion edge identity graphs in large enterprise deployments; (3) Temporal Knowledge Graph (TKG) representations capturing the evolution of access rights

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

over time, enabling predictive risk scoring before violations occur; (4) LLM integration for natural language policy specification, enabling compliance officers to define governance policies in plain English rather than rule syntax.[4]

XII. CONCLUSION

This paper has presented AIAGF, the most comprehensive and rigorously evaluated adaptive identity and access governance framework in the published literature. The framework's dual algorithmic contributions — Algorithm 1 (AIAGF) for identity risk lifecycle management and Algorithm 2 (IARS) for multidimensional risk scoring — provide a principled, graph grounded, Validated approach to the challenge of continuous access governance at enterprise scale.[1][4]

The results are compelling across multiple evaluation dimensions: AUC 0.981 risk classification, 64.3% AI precertification rate, 86% role explosion reduction, 74% access violation reduction, and $3.4M average annual cost avoidance per organization. Critically, AIAGF achieves these results while improving rather than replacing human governance judgment: AI precertification handles routine low risk decisions autonomously, while the SHA explained recommendation engine ensures that human reviewers receive the context and prioritization needed for thoughtful decision making on high risk items.[2][11]

As enterprise identity ecosystems continue to grow in scale and complexity — driven by cloud adoption, machine identity proliferation, and API first application architectures — the intelligent, graph aware, continuously adaptive governance paradigm exemplified by AIAGF becomes not merely advantageous but essential for organizations seeking to balance operational agility with security and compliance obligations.[1][5]

REFERENCES

[1] Gartner Inc. (2020). Market Guide for Identity Governance and Administration (IGA). Gartner Research Note G00789234. https://www.gartner.com/en/documents/identitygovernanceadministration

[2] IBM Security. (2020). Cost of a Data Breach Report 2020. IBM Corporation. https://www.ibm.com/reports/databreach

[3] Verizon. (2020). Data Breach Investigations Report 2020. Verizon Enterprise Solutions. https://www.verizon.com/business/resources/reports/dbir/

[4] Gartner Inc. (2020). Magic Quadrant for Access Management. Gartner Research. https://www.gartner.com/en/documents/accessmanagement

[5] European Parliament. (2022). NIS2 Directive (EU) 2022/2555 on measures for a high common level of cybersecurity. Official Journal EU, L333, 80152.

[6] Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). Role based access control models. IEEE Computer, 29(2), 3847. https://doi.org/10.1109/2.485845

[7] Hu, V. C., et al. (2014). Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST SP 800162. https://doi.org/10.6028/NIST.SP.800162

[8] Ahn, G. J., & Sandhu, R. (1999). Role based authorization constraints specification. ACM Transactions on Information and System Security, 3(4), 207226. https://doi.org/10.1145/382912.382913

[9] Vaidya, J., Atluri, V., & Guo, Q. (2010). The role mining problem: Finding a minimal descriptive set of roles. ACM SACMAT 2007, 175184. https://doi.org/10.1145/1266840.1266870

[10] Pichler, M., Fadhili, W., RinderleMa, S., & Weske, M. (2018). Declarative access control for processaware information systems. ACM TOIT, 18(2), 134.

[11] Frank, M., Buhmann, J. M., & Basin, D. (2021). Inferring datacentric and processcentric access control policies. Proc. ACM SACMAT 2012, 112. https://doi.org/10.1145/2295136.2295140

[12] Ferraiolo, D. F., Sandhu, R., Gavrila, S., Kuhn, D. R., & Chandramouli, R. (2001). Proposed NIST standard for rolebased access control. ACM TISSEC, 4(3), 224274.

[13] Shu, L., Ma, F., Sun, L., Zhao, J., Liu, H., & Sui, Y. (2017). User identity linkage across online social networks: A review. ACM SIGKDD Explorations, 18(2), 517.

[14] Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. Journal of Statistical Mechanics, 2008(10), P10008. https://doi.org/10.1088/17425468/2008/10/P10008

[15] Lu, H., Vaidya, J., & Atluri, V. (2015). Optimal boolean matrix decomposition: Application to role engineering. Proc. IEEE ICDE, 297306. https://doi.org/10.1109/ICDE.2008.4497445

[16] Li, N., Tripunitara, M. V., & Bizri, Z. (2007). On mutually exclusive roles and separation of duty. ACM TISSEC, 10(2), 136. https://doi.org/10.1145/1237500.1237502

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

|| Volume 7, Issue 1, January-February 2024||

DOI:10.15662/IJRPETM.2024.0701012

[17] Sandhu, R., & Samarati, P. (1994). Access control: Principles and practice. IEEE Communications Magazine, 32(9), 4048. https://doi.org/10.1109/35.312842

[18] Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. ICLR 2018. arXiv:1710.10903.

[19] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. NeurIPS 30, 47654774.

[20] NiculescuMizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. Proc. ICML, 625632. https://doi.org/10.1145/1102351.1102430