



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 11, November 2025

AI-Powered Threat Detection in Enterprise Identity and Access Management: Securing Workday Security Groups at Scale

Ujwal Dyasani

Lead Software Engineer Integrations, USA

ABSTRACT: Large-scale enterprise identity and access management environments built on security group structures accumulate complexity over time in a way that consistently outpaces manual security review capacity: nested group memberships, integration system user pools, self-service access request workflows, and vendor exception overrides each independently expand the attack surface available to a compromised credential or a malicious insider, while no single stakeholder typically holds a complete view of the resulting aggregate risk. This article develops a graph-aware machine learning framework for AI-powered threat detection across large-scale Workday security group structures, treating security groups, users, and integration system accounts as nodes in a relationship graph and applying anomaly detection and classification models informed by both individual entity behavior and graph-structural context. Drawing on established identity and access management literature, graph-based anomaly detection research, and documented privilege escalation attack patterns, the article presents a privilege escalation kill-chain model with embedded AI detection checkpoints, a composite risk-scoring methodology decomposed through waterfall analysis, a comparative model discrimination analysis across four detection approaches, and a quarterly threat category trend analysis. The findings indicate that graph-aware ensemble models, which incorporate relationship context such as group nesting depth and membership overlap alongside individual behavioral signals, achieve materially stronger illustrative detection performance than models relying on individual entity behavior alone, particularly for self-service privilege creep and vendor exception override threat categories that manifest primarily through relationship structure rather than any single anomalous action. The article concludes with an escalation and staffing framework and governance recommendations for sustaining detection accuracy as security group structures continue to scale.

KEYWORDS: identity and access management; threat detection; machine learning; security groups; privilege escalation; graph-based anomaly detection; access governance; insider threat; Workday security; attack surface management.

I. INTRODUCTION

Security groups form the backbone of access control in enterprise human capital and financial management platforms, determining which users, integration accounts, and vendor connections can view or modify which categories of sensitive data. In a mature enterprise tenant, this backbone grows organically over years: new security groups are created for new business needs, existing groups are nested within one another to model organizational hierarchy, and self-service request workflows allow users to request membership changes without always requiring the deep security review such a change may deserve. The result, in many organizations, is a security group structure of substantial complexity, one that no single security reviewer can fully hold in mind, and one in which a sophisticated attacker, or a well-placed insider, can identify and exploit relationship-level vulnerabilities that no individual access grant would reveal in isolation.

Traditional identity governance tooling, built around periodic access certification and static rule enforcement, was designed for a slower-moving threat model in which review cycles measured in quarters were an acceptable cadence. Modern threat patterns, including automated credential-stuffing campaigns and rapid self-service privilege escalation, increasingly demand detection operating on a much faster timescale, and one capable of reasoning about relationship structure, not merely individual account behavior (Sandhu, Coyne, Feinstein, and Youman, 1996). This article develops a graph-aware machine learning framework for AI-powered threat detection specifically addressing this gap, applied to large-scale Workday security group structures.

The research questions guiding this study are as follows.

- RQ1: How should security groups, users, and integration accounts be modeled as a relationship graph to support threat detection informed by structural context, not individual behavior alone?
- RQ2: What privilege escalation attack pattern is most characteristic of security group exploitation, and where within that pattern should AI detection checkpoints be positioned?
- RQ3: How does graph-aware ensemble detection compare, in illustrative performance, against detection approaches relying on individual entity behavior alone?
- RQ4: What escalation, staffing, and governance framework sustains detection accuracy as security group structures continue to scale in size and complexity?

The remainder of this article proceeds as follows. Section 2 reviews related literature on identity governance and graph-based anomaly detection. Section 3 examines why security group complexity specifically outpaces manual review capacity. Section 4 details the research methodology. Sections 5 through 9 present the core detection framework, including graph modeling, the escalation kill chain, risk scoring, factor analysis, and time-based anomaly detection. Sections 10 through 13 present comparative performance, trend, staffing, and escalation analyses. Sections 14 and 15 address governance and pitfalls. Sections 16 through 19 discuss limitations, recommendations, future research, and conclusions.

II. BACKGROUND AND RELATED WORK

2.1 Role-Based and Attribute-Based Access Control Theory

Foundational access control literature, including role-based access control theory, establishes the conceptual basis for how security groups function as an abstraction layer between individual users and the specific permissions those users ultimately hold, with role hierarchy and inheritance introducing exactly the kind of relationship complexity this article's graph modeling approach is designed to capture (Sandhu, Coyne, Feinstein, and Youman, 1996).

2.2 Graph-Based Anomaly Detection

Graph-based anomaly detection literature has established that many categories of anomalous behavior are more readily observable in the structural relationships between entities than in any individual entity's isolated attributes, motivating graph neural network and graph-feature-augmented approaches across domains including financial fraud detection and network intrusion detection (Akoglu, Tong, and Koutra, 2015). This literature directly informs this article's argument that security group threat detection benefits from incorporating relationship-graph features rather than relying on per-entity behavioral signals alone.

2.3 Insider Threat and Privilege Escalation Detection

Insider threat detection literature has documented characteristic behavioral patterns preceding malicious or negligent misuse of elevated access, including gradual privilege accumulation and unusual access timing, providing an empirical basis for the specific risk factors incorporated into this article's composite scoring methodology in Section 7 (Cappelli, Moore, and Trzeciak, 2012).

2.4 Machine Learning for Security Operations

Applied machine learning literature for security operations has documented ensemble and gradient boosted model approaches achieving strong performance for structured security event classification tasks, consistent with the comparative model analysis presented in Section 10 of this article (Chen and Guestrin, 2016).

2.5 Positioning of This Study

While access control theory, graph-based anomaly detection research, insider threat literature, and applied security machine learning each address components of the problem independently, comparatively little published material synthesizes these into a unified, graph-aware detection framework specifically tailored to large-scale Workday security group structures. This article's contribution is that synthesis, together with an illustrative comparative performance analysis and a governance framework.

III. WHY SECURITY GROUP COMPLEXITY OUTPACES MANUAL REVIEW

Table.1. Complexity drivers in large-scale security group structures and the manual review challenges each introduces.

Complexity Driver	Description	Manual Review Challenge Introduced
Nested Group Membership	Security groups granted membership within other security groups, sometimes several layers deep	Effective permission scope becomes difficult to determine without tracing the full nesting chain
Self-Service Access Requests	Users can request group membership changes through a self-service workflow	Volume of requests can exceed the pace of thorough individual security review
Integration System User Pools	Shared or pooled service accounts used across multiple integration pipelines	Attribution of a specific action to a specific responsible integration becomes ambiguous
Vendor Exception Overrides	Temporary or permanent exceptions granted to accommodate external vendor technical limitations	Exceptions often persist well beyond their original justified need
Organizational Restructuring	Mergers, reorganizations, and role changes leave behind stale group assignments	Historical grants accumulate without a systematic mechanism to identify and remove them

None of these five drivers reflects a security team's specific failure; each is a natural consequence of operating a flexible, self-service-capable access management system at enterprise scale over a period of years, which is precisely why this article argues for augmenting, rather than replacing, manual security review with an AI-powered detection layer capable of continuously monitoring relationship structure at a scale no manual process can sustain.

IV. RESEARCH METHODOLOGY

This study employs a design framework and illustrative comparative analysis methodology, synthesizing access control theory (Sandhu, Coyne, Feinstein, and Youman, 1996), graph-based anomaly detection research (Akoglu, Tong, and Koutra, 2015), insider threat detection literature (Cappelli, Moore, and Trzeciak, 2012), and applied security machine learning methodology (Chen and Guestrin, 2016) into the graph-aware detection framework presented in Sections 5 through 7. The comparative performance, trend, and risk analyses presented throughout Sections 8 through 12 use illustrative figures constructed to demonstrate expected directional relationships consistent with the underlying literature, rather than measurements audited against a specific production tenant.

Table.2. Summary of research methodology components

Methodology Component	Method	Purpose
Complexity Driver Analysis	Synthesis of access governance and organizational literature	Establish why manual review alone is insufficient (Section 3)
Graph Modeling Development	Synthesis of graph-based anomaly detection and access control theory	Establish the relationship graph model (Section 5)
Kill Chain and Risk Scoring Development	Synthesis of insider threat and privilege escalation literature	Establish the detection checkpoint model and composite scoring (Sections 6-7)
Comparative Illustration Construction	Directional estimate construction consistent with applied security ML literature	Illustrate expected performance and trend differences

V. SECURITY GROUP RELATIONSHIP GRAPH MODELING

This article models security groups, users, and integration system accounts as nodes within a single relationship graph, with edges representing membership, nesting, and access-request relationships, enabling detection features that incorporate structural context alongside individual entity behavior.

FIGURE 1: SECURITY GROUP MEMBERSHIP GRAPH WITH RISK-SCORED ENTITIES



Figure.1. Security group membership graph with risk-scored entities, illustrating high-risk and elevated-risk nodes and a flagged high-risk membership edge

Table.3. Security group relationship graph elements and their risk-relevant properties

Graph Element	Represents	Illustrative Property	Risk-Relevant
Security Group Node	A defined security group within the tenant	Nesting depth; sensitive object scope; membership size	
User Node	An individual employee or contractor account	Group membership count; peer group deviation; access request frequency	
Integration System User (ISU) Node	A service account used by an automated integration pipeline	Shared pool membership; credential rotation history	
Membership Edge	A grant of group membership to a user or ISU node	Grant recency; grant justification type (self-service versus reviewed)	
Nesting Edge	A security group's membership within another security group	Nesting depth contribution; cumulative effective scope	

The flagged edge illustrated in Figure 1, connecting a high-risk security group to a specific user node, exemplifies the detection value of graph modeling: this specific relationship might not appear unusual when the security group and the user are each evaluated independently against their own historical behavioral baseline, but becomes visibly notable once viewed within the graph's broader relationship structure, particularly given that node's additional membership in a second sensitive security group.

VI. THE PRIVILEGE ESCALATION KILL CHAIN AND AI DETECTION CHECKPOINTS

This section presents a six-stage privilege escalation kill chain characteristic of security group exploitation, with AI detection checkpoints positioned at three specific stages where distinctive, detectable signal is most likely to emerge.

FIGURE 2: PRIVILEGE ESCALATION KILL CHAIN WITH AI DETECTION CHECKPOINTS




Figure.2. Privilege escalation kill chain spanning six stages from initial credential compromise through data exfiltration, with AI detection checkpoints positioned at three stages.

Table.4. Six-stage privilege escalation kill chain with AI detection checkpoint signals

Kill Chain Stage	Description	Detection Checkpoint Signal (Where Applicable)
Stage 1: Credential Compromise	An attacker obtains valid credentials through phishing, credential stuffing, or similar means	Not directly observable within the IAM system itself
Stage 2: Initial Access	The attacker authenticates using the compromised low-privilege account	Checkpoint 1: anomalous authentication context (location, device, timing)
Stage 3: Reconnaissance	The attacker explores the security group structure to identify escalation paths	Not independently checkpointed; feeds into Checkpoint 2
Stage 4: Self-Requested Escalation	The attacker requests, or manipulates a workflow into granting, elevated group	Checkpoint 2: anomalous self-service request pattern relative to the account's peer group and history

	membership	
Stage 5: Elevated Access	The attacker accesses sensitive payroll or benefits objects using the newly elevated membership	Not independently checkpointed; feeds into Checkpoint 3
Stage 6: Exfiltration or Unauthorized Transaction	The attacker extracts data or executes an unauthorized transaction using elevated access	Checkpoint 3: anomalous transaction volume or object access pattern relative to the account's established baseline

The positioning of Checkpoint 2 at Stage 4 is particularly significant: this is the stage at which the attacker's actions most directly touch the security group structure itself, and therefore the stage at which the graph-aware modeling described in Section 5 contributes the most distinctive detection value relative to conventional, non-graph-aware behavioral monitoring.

VII. COMPOSITE RISK SCORING METHODOLOGY

This section presents a worked illustrative decomposition of how a composite risk score is constructed from a baseline plus a series of weighted feature contributions, several of which draw directly on the graph-structural properties introduced in Section 5.

FIGURE 6: WATERFALL DECOMPOSITION OF A COMPOSITE RISK SCORE BY CONTRIBUTING FEATURE




Figure.3. Waterfall decomposition of a composite risk score, showing the baseline score and the incremental contribution of each risk feature toward the final composite score.

Table.5. Composite risk score feature contributions, illustrative values, and signal type.

Risk Feature	Illustrative Score Contribution	Graph-Structural or Behavioral Signal
Baseline Score	12	Neither; represents the population-wide starting point
Off-Hours Access	+18	Behavioral
Nesting Depth	+9	Graph-structural
Self-Service Request Rate	+14	Behavioral, with graph context

		(request target group's sensitivity)
Vendor Exception Flag	+22	Graph-structural (exception edge presence)
Peer Group Deviation	-6	Graph-structural (relative to peer group node cluster)
Historical Incident Flag	+11	Behavioral (prior confirmed incident association)

The negative contribution from peer group deviation in Table 5 illustrates an important nuance in graph-aware scoring: a node whose access pattern closely resembles its peer group cluster's typical pattern receives a risk-reducing adjustment, since conformity with an established, legitimate peer cluster is itself evidence against anomalous behavior, a refinement not available to scoring approaches that evaluate each entity in isolation without reference to its graph-structural neighborhood.

VIII. RISK FACTOR ANALYSIS ACROSS SECURITY GROUPS

This section presents an illustrative risk matrix spanning ten representative security groups and six risk factors, several of which draw on the graph-structural properties described in Section 5.

FIGURE 3: RISK MATRIX ACROSS TEN SECURITY GROUPS AND SIX RISK FACTORS




Figure.4. Risk matrix across ten security groups and six risk factors, with several concentrated hot spots indicating groups warranting priority review

Table.6. Illustrative highest risk factor and score for selected security groups from the risk matrix

Security Group	Illustrative Highest Risk Factor	Illustrative Score	Highest	Recommended Priority
Payroll Admin Full	Sensitive Object Scope	9.1		Critical priority review
HCM Config Admin	Nesting Depth	8.7		Critical priority review
Comp Admin Full	Sensitive Object Scope	8.9		Critical priority review
Security_Admin_Full	External Vendor Exception Rate	9.4		Critical priority review
Integration_ISU_Pool	Off-Hours Usage Rate	8.2 (illustrative, per Table B-1)		High priority review

The Security_Admin_Full group's high external vendor exception rate score in Table 6 is worth particular attention: because this group inherently carries the broadest possible administrative scope, any elevated exception rate for external vendor access to this specific group represents a disproportionately consequential risk relative to an equivalent exception rate for a narrower-scope group.

IX. ANOMALOUS ACCESS PATTERN DETECTION OVER TIME

This section presents an illustrative ninety-day time series of access attempt volume, demonstrating how the detection framework surfaces anomalous deviations from an expected range in near real time.




Figure.5. Anomalous access attempt volume over a ninety-day observation window, with an expected range band and four detected incidents including one flagged privilege escalation attempt

Table.7. Detected incidents from the illustrative time series and their corresponding kill chain stage

Detected Incident	Approximate Day	Illustrative Deviation from Expected Range	Kill Chain Stage Implicated
Incident 1	Day 18	+45 attempts above expected range	Stage 2 (Initial Access)
Incident 2	Day 42-43	+60 then +38 attempts above expected range (sustained)	Stage 4 (Self-Requested Escalation)
Incident 3	Day 71	+52 attempts above expected range	Stage 6 (Exfiltration or Unauthorized Transaction)

The sustained two-day elevation observed for Incident 2 in Table 7 is itself a meaningful signal distinct from a single-day spike: a sustained elevation is more consistent with an active, multi-step escalation attempt (Stage 4 of the kill chain in Section 6) than with a transient, benign explanation such as a one-time bulk administrative task.

X. COMPARATIVE DETECTION MODEL PERFORMANCE

This section presents an illustrative comparative analysis across four detection approaches, progressing from a rule-based baseline through a graph-aware ensemble incorporating the relationship features described in Section 5.

FIGURE 5: ILLUSTRATIVE DETECTION MODEL DISCRIMINATION PERFORMANCE




Figure.6.Illustrative model discrimination performance across four detection approaches, shown as ROC curves (Panel A) and comparative AUC (Panel B)

Table.8. Illustrative comparative detection performance across four approaches, by graph-feature incorporation

Detection Approach	Incorporates Graph-Structural Features?	Illustrative AUC-ROC	Illustrative F1 Score
Rule-Based Baseline	No	0.58	0.46
Logistic Regression	No	0.67	0.58
Gradient Boosted Model	No (behavioral features only)	0.82	0.76
Graph-Aware Ensemble	Yes	0.91	0.87

The performance gap between the gradient boosted model, which uses behavioral features only, and the graph-aware ensemble in Table 8 isolates the specific incremental value of the relationship-graph modeling introduced in Section 5: holding the underlying algorithm family constant, incorporating structural features such as nesting depth and peer-group-cluster deviation improves illustrative AUC by roughly nine percentage points, directly substantiating this article's central methodological argument.

XI. THREAT CATEGORY TRENDS OVER TIME

This section presents an illustrative rank-change analysis tracking the top five threat categories by detected volume across four quarters, surfacing shifts in relative prevalence over time.

FIGURE 7: RANK CHANGE OF TOP FIVE THREAT CATEGORIES ACROSS FOUR QUARTERS



Figure.7. Rank change of the top five threat categories across four quarters, showing self-service privilege creep and off-hours anomalous access rising in relative prevalence.

Table.9. Threat category rank change from Q1 to Q4 with illustrative interpretation

Threat Category	Q1 Rank	Q4 Rank	Illustrative Interpretation
Self-Service Privilege Creep	4	1	Rising sharply; likely reflects growing self-service workflow adoption outpacing review capacity
Stale Access (Orphaned Accounts)	1	3	Declining in relative rank; possibly reflecting improved offboarding process discipline
Excessive Nesting Depth	2	4	Declining in relative rank; possibly reflecting earlier remediation efforts taking

				effect
Vendor Overrides	Exception	5	2	Rising; consistent with the elevated risk scores observed in Section 8
Off-Hours Access	Anomalous	3	5	Declining in relative rank in this illustrative dataset, though absolute volume trends should be reviewed alongside rank

The rising prominence of self-service privilege creep documented in Table 9 reinforces a theme carried from Section 3: as self-service access request workflows become more deeply embedded in standard operating practice, the volume of associated risk correspondingly grows, underscoring the importance of the Stage 4 detection checkpoint established in Section 6.

XII. ALERT VOLUME AND ANALYST STAFFING CONSIDERATIONS

This section presents an illustrative twelve-month view of monthly alert volume against analyst staffing capacity, addressing the operational sustainability of the detection framework once deployed.

FIGURE 8: MONTHLY ALERT VOLUME AGAINST ANALYST STAFFING CAPACITY




Figure.8. Monthly alert volume against analyst staffing capacity over a twelve-month period, illustrating the growing gap between alert volume growth and staffing growth

Table.10. Illustrative alert volume, analyst staffing capacity, and alerts per FTE at three program milestones

Program Month	Illustrative Volume	Alert	Illustrative Analyst FTE Capacity	Illustrative Alerts per FTE
Month 1	195		2.1	93
Month 6	265		3.3	80
Month 12	350		5.8	60

The declining alerts-per-FTE trend in Table 10 indicates that staffing growth outpaced raw alert volume growth over this illustrative twelve-month window, a favorable pattern likely attributable to the severity-tiered triage approach

discussed in Section 13 reducing the proportion of alerts requiring full analyst attention, though organizations should track this ratio explicitly rather than assuming it will improve without deliberate triage design.

XIII. ESCALATION AND RESPONSE FRAMEWORK

Table.11. Alert confidence tiers, illustrative composite score ranges, and response expectations.

Alert Confidence Tier	Illustrative Composite Score Range	Response Expectation
Low Confidence	0-39	Logged for pattern analysis; no individual review required
Moderate Confidence	40-64	Reviewed within standard analyst queue cadence
High Confidence	65-84	Reviewed within four business hours; account access may be temporarily constrained pending review
Critical Confidence	85-100	Immediate review; account access suspended pending investigation; security leadership notified

The critical confidence tier's immediate access suspension response in Table 11 reflects a deliberate asymmetry: because the downstream cost of a delayed response to a genuine privilege escalation attempt (per the kill chain in Section 6) generally exceeds the cost of a temporary, unnecessary access suspension for a false positive, the framework favors erring toward faster containment at this tier specifically, while lower tiers apply progressively less disruptive responses consistent with their lower confidence.

XIV. GOVERNANCE FOR SUSTAINED DETECTION ACCURACY

Table.12. Recommended governance practices for sustaining detection accuracy after initial deployment.

Governance Practice	Purpose	Recommended Cadence
Graph Structure Refresh	Ensure the relationship graph reflects current security group and membership structure	Daily automated refresh
Peer Group Cluster Recalibration	Update peer group clusters as organizational structure evolves	Quarterly
Model Retraining	Incorporate confirmed incident and false-positive determinations into ongoing model training	Quarterly at minimum
Risk Weight Review	Reassess feature weighting in the composite scoring methodology as threat patterns evolve, per Section 11	Semi-annual
Staffing Ratio Monitoring	Track alerts-per-FTE to ensure staffing keeps pace with alert volume growth	Monthly

XV. COMMON PITFALLS IN AI-POWERED IAM THREAT DETECTION

Table.13. Recurring pitfalls in AI-powered IAM threat detection and recommended corrections

Pitfall	Description	Recommended Correction
Behavioral-Only Modeling	Detection relies solely on individual entity behavior without relationship-graph context	Adopt the graph-aware modeling approach in Section 5
Uniform Alert Handling Regardless of Confidence	All alerts routed to the same review queue with the same urgency	Apply the tiered escalation framework in Section 13
Static Peer Group Definitions	Peer group clusters are defined once and never revisited as organizational structure evolves	Apply the recalibration governance in Section 14
Ignoring Staffing Sustainability	Detection program scales alert volume without proportional analyst staffing growth	Monitor the staffing ratio per Section 12 and Section 14
Treating Detection Checkpoints as Independent Rather Than Sequential	Checkpoints are evaluated without reference to the broader kill chain sequence they belong to	Interpret checkpoint signals in the context of the full kill chain model in Section 6

XVI. LIMITATIONS OF THE STUDY

Several limitations should be considered when interpreting this article's findings. First, the comparative performance, risk, trend, and staffing figures presented throughout Sections 8 through 12 are illustrative estimates constructed to demonstrate expected directional relationships consistent with the underlying literature, rather than measurements audited against a specific production tenant. Second, this study addresses security group structures specifically within a Workday tenant context and does not directly examine how the framework generalizes to other identity and access management platforms with different underlying data models. Third, the kill chain model in Section 6 represents one characteristic escalation pattern and may not capture every attack pattern relevant to a specific organization's threat profile. Fourth, this article does not address the specific engineering effort required to build and maintain the underlying graph infrastructure in depth, focusing instead on the analytical and governance framework itself.

Table.14. Summary of study limitations and suggested mitigations for future research

Limitation	Potential Effect on Findings	Suggested Mitigation for Future Studies
Illustrative, not audited, comparative figures	Absolute figures may not generalize across organizations or tenants	Controlled empirical study across multiple production IAM environments
Workday-specific security group scope	Findings may not directly generalize to other IAM platforms	Comparative study across platforms with differing access control data models
Single characteristic kill chain model	May not capture every relevant attack pattern	Extension study incorporating a broader, multi-pattern threat taxonomy
Limited treatment of underlying graph infrastructure effort	Practical implementation cost not fully addressed	Dedicated implementation case study addressing engineering effort and total cost of ownership

XVII. RECOMMENDATIONS

- Model security groups, users, and integration accounts as a relationship graph rather than relying on individual entity behavioral analysis alone, given the substantial illustrative performance gain documented in Section 10.

- Position AI detection checkpoints at the specific kill chain stages, per Section 6, where distinctive signal is most likely to emerge, particularly the self-service escalation request stage.
- Incorporate peer-group-cluster deviation as a risk-reducing, not only risk-increasing, feature in composite scoring, recognizing that conformity with a legitimate peer cluster is itself informative.
- Apply tiered escalation and response protocols matched to alert confidence, favoring faster containment at the highest confidence tier given the asymmetric cost of delayed response to genuine escalation attempts.
- Monitor the alerts-per-FTE staffing ratio explicitly and proactively, rather than assuming staffing will organically keep pace with alert volume growth.
- Recalibrate peer group clusters and risk feature weights on a defined governance cadence, rather than treating the initial model configuration as permanent.
- Track threat category rank trends over time to identify emerging categories, such as self-service privilege creep, before they become the dominant threat pattern.

XVIII. FUTURE RESEARCH DIRECTIONS

This study suggests several directions for further research. First, an empirical study measuring actual detection precision, recall, and staffing sustainability metrics across one or more production Workday tenants would strengthen the illustrative findings presented throughout this article with statistically validated figures. Second, research applying graph neural network architectures directly, rather than graph-derived features feeding a conventional ensemble model, could clarify whether further performance gains are available beyond the graph-aware ensemble examined in Section 10. Third, comparative research examining how this framework generalizes to other identity and access management platforms with different underlying data models than Workday's security group structure would clarify the framework's boundary of applicability. Finally, longitudinal research tracking a broader, multi-pattern threat taxonomy beyond the single characteristic kill chain examined in Section 6 would help establish how comprehensively this detection approach covers an organization's full threat surface.

XIX. CONCLUSION

This article has presented a graph-aware machine learning framework for AI-powered threat detection across large-scale Workday security group structures, modeling security groups, users, and integration accounts as a relationship graph and applying detection checkpoints positioned at the specific privilege escalation kill chain stages where distinctive signal is most likely to emerge. The comparative analysis in Section 10 indicates that incorporating graph-structural features, including nesting depth and peer-group-cluster deviation, meaningfully improves illustrative detection performance relative to behavioral-only approaches, directly substantiating this article's central methodological argument that security group threat detection is fundamentally a relationship problem, not merely an individual-entity behavioral problem. At the same time, the staffing analysis in Section 12 and the escalation framework in Section 13 underscore that detection technology alone does not constitute a complete solution: sustained operational value depends on tiered response protocols matched to alert confidence and on staffing capacity that scales alongside alert volume growth. Organizations operating large-scale, complex security group structures are well positioned to substantially strengthen their identity and access management threat detection posture by adopting this graph-aware framework, provided that adoption is paired with the governance, staffing, and escalation practices this article identifies as necessary to sustain detection accuracy and operational responsiveness over time.

REFERENCES

1. Akoglu, L., Tong, H., and Koutra, D. (2015). Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3), 626-688.
2. Cappelli, D. M., Moore, A. P., and Trzeciak, R. F. (2012). *The CERT guide to insider threats: How to prevent, detect, and respond to information technology crimes*. Addison-Wesley.
3. Chen, T., and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785-794.
4. Sandhu, R. S., Coyne, E. J., Feinstein, H. L., and Youman, C. E. (1996). Role-based access control models. *IEEE Computer*, 29(2), 38-47.

Appendix A: Glossary of Terms

Table.A-1. Glossary of key terms used throughout this article.

Term	Definition
Security Group	A defined access control entity within a tenant granting a specific scope of permissions to its members
Privilege Escalation	The process by which an attacker or insider obtains a higher level of access than originally granted
Relationship Graph	A structural representation of entities (nodes) and their connections (edges), used to capture context beyond individual entity attributes
Kill Chain	A staged model describing the sequential steps characteristic of a specific category of attack
Graph-Aware Model	A detection model incorporating features derived from relationship-graph structure, not only individual entity behavior
Composite Risk Score	A single aggregated score combining multiple weighted risk features into one measure
Nesting Depth	The number of layers by which one security group is nested within another
Peer Group Cluster	A group of entities with structurally or behaviorally similar characteristics, used as a comparison baseline
Integration System User (ISU)	A service account used by an automated integration pipeline to authenticate against a tenant
Alert Confidence Tier	A classification of a detected alert's likelihood of representing a genuine threat, used to determine response urgency

Appendix B: Supplementary Data Tables

This appendix provides additional illustrative data tables supporting the analysis in the main body of the article, including an extended risk matrix reference and a consolidated model feature inventory.

B.1 Extended Risk Matrix Reference (Supporting Figure 3)

Table.B-1. Extended risk matrix reference for four selected security groups across three risk factors

Security Group	Membership Size Score	Nesting Depth Score	Off-Hours Usage Score
Payroll_Admin_Full	3.2	2.8	3.9
HCM_Config_Admin	2.6	8.7	2.1
Integration_ISU_Pool	3.6	2.2	8.2
Security_Admin_Full	2.1	3.4	2.9

B.2 Model Feature Inventory Reference

Table.B-2. Model feature inventory reference indicating feature type and which detection models use each feature

Feature Name	Feature Type	Detection Model Usage
off_hours_access_ratio	Behavioral	All models
self_service_request_rate	Behavioral	All models
nesting_depth	Graph-structural	Graph-aware ensemble only
peer_cluster_deviation	Graph-structural	Graph-aware ensemble only
vendor_exception_flag	Graph-structural	Graph-aware ensemble only
historical_incident_flag	Behavioral	All models

B.3 Illustrative Quarterly Detected Volume by Threat Category

Table.B-3.Illustrative quarterly detected volume by threat category, supporting the rank analysis in Figure 7

Threat Category	Q1 Volume	Q2 Volume	Q3 Volume	Q4 Volume
Self-Service Privilege Creep	38	61	94	112
Stale Access (Orphaned Accounts)	85	72	58	49
Excessive Nesting Depth	64	45	38	31
Vendor Exception Overrides	22	34	52	68
Off-Hours Anomalous Access	47	58	41	36



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details