

A Theoretical Risk-Scoring Model for Cloud IAM Complexity

MSc Cybersecurity & Digital Forensics
Research Methods & Professional Development

Research Proposal

Shamsuddin Mahmud
Student ID: 2432880

Abstract

Cloud Identity and Access Management (IAM) is a primary control plane for authorisation in public cloud platforms, yet IAM configurations can grow rapidly in size and interdependence. This proposal outlines a theoretical, graph-informed risk-scoring model to quantify IAM complexity and translate structural complexity signals into a defensible “IAM complexity risk score”. The study will identify measurable complexity factors, define a weighted scoring approach, represent IAM entities/relationships as a graph, and conceptually validate the model against hypothetical IAM scenarios. The intended contribution is an auditable, repeatable way to reason about IAM complexity as a risk driver, supporting more proactive cloud security governance.

Contents

Abstract.....	1
1. Introduction.....	3

1.1 Context and Background.....	3
1.2 Research Problem	3
1.3 Research aim and objectives	4
1.4 Significance of the research	4
2. Literature review	4
2.1 Cloud IAM misconfiguration and excessive privilege	5
2.2 IAM complexity drivers (privilege creep, policy sprawl, role explosion).....	5
2.3 Graph-based analysis, trust relationships, and attack paths	5
2.4 Risk assessment frameworks (NIST SP 800-30 and FAIR).....	6
2.5 Literature summary and gap statement	6
3. Research Design and Methodology	7
3.1 Research design	7
3.2 Research methods and materials	7
3.2.1 Methods.....	7
3.2.2 Materials	8
4. Expected outcome	9
5. Timeline	9
7. References	10

1. Introduction

1.1 Context and Background

Cloud computing adoption continues to accelerate across sectors, increasing reliance on shared, programmable infrastructure and managed platform services. In these environments, Identity and Access Management (IAM) becomes the central mechanism that authenticates principals (human and non-human identities) and authorises actions against cloud resources. As a result, IAM configuration effectively defines an organisation's security perimeter: every API call is mediated by an identity and evaluated against policy [3].

However, IAM deployments rarely remain simple. Modern cloud estates feature rapid provisioning, microservices, CI/CD automation, and cross-account or cross-tenant integrations, all of which increase the number of roles, policies, groups, service accounts, and trust relationships. Over time, operational pressures can drive "privilege creep" (permissions accumulate beyond what is needed), "policy sprawl" (policies proliferate and become hard to reason about), and "role explosion" (many fine-grained roles created with overlapping privileges). These patterns are widely recognised as risk multipliers because they can hide unintended privilege escalation paths and expand attack surface [6], [10].

Recent practitioner and research trends increasingly emphasise relationship-aware cloud security: rather than treating findings (e.g., a single excessive permission) in isolation, organisations use contextual modelling to understand how misconfigurations combine into end-to-end attack paths. Graph-based approaches are particularly well suited for this purpose because they represent identities, resources, and permissions as nodes and edges, enabling path analysis and prioritisation [6], [12]. Tools and discussions around cloud attack path analysis and graph modelling illustrate the growing recognition that "connectivity" and "reachability" matter as much as the presence of individual vulnerabilities [6].

1.2 Research Problem

Problem statement: Cloud IAM configurations can reach a level of structural complexity where security teams cannot reliably reason about effective permissions, transitive trust relationships, and privilege escalation paths. Although IAM misconfiguration is frequently linked to breaches and unauthorised access, organisations lack a standard, defensible method to *quantify* IAM complexity as a risk driver and use that measurement to assess risk proactively [7], [10].

Current challenges and gaps: Current IAM assurance practices often depend on periodic access reviews, policy linting, and detection-oriented tools that enumerate misconfigurations or "risky permissions" but do not provide a unified measure of overall IAM complexity [10]. Risk frameworks such as NIST SP 800-30 provide structured guidance for conducting risk assessments, but they are not IAM-specific and do not directly operationalise IAM structural complexity (e.g., density of trust relationships, role chaining, or cumulative privilege growth)

into a scoring model [1]. Similarly, quantitative risk approaches such as FAIR decompose risk into frequency and magnitude factors, but they do not prescribe how to derive those factors from cloud IAM graph structure [2]. Consequently, organisations struggle to compare IAM estates over time, set measurable governance thresholds, or prioritise remediation based on the structural “riskiness” of the identity-permission network [6].

1.3 Research aim and objectives

Aim: To develop a theoretical risk-scoring model that quantifies IAM complexity in cloud environments.

- Identify measurable IAM complexity factors such as privilege creep, policy sprawl, trust relationship density, least-privilege deviation, attack path complexity, and role explosion.
- Develop a weighted scoring model that combines these factors into an overall IAM complexity risk score.
- Represent IAM entities and relationships using a graph-based structure suitable for path analysis and complexity measurement.
- Validate the model conceptually using hypothetical IAM scenarios and sensitivity testing of weights and factor definitions.

1.4 Significance of the research

Importance: IAM is routinely described as the control plane for cloud security; therefore, unmanaged complexity in IAM can translate into systemic exposure (hidden privilege escalation routes and wider blast radius if a principal is compromised) [6], [10]. A theoretical scoring model for IAM complexity can contribute to the research community by formalising how structural characteristics of IAM (relationships, reachability, and permission aggregation) relate to risk concepts such as likelihood and impact [1], [2]. This work also responds to a gap noted in the literature and practice: the lack of an IAM-specific quantification approach for complexity-driven risk [6].

Potential applications: In practice, an IAM complexity risk score could be used to (i) benchmark IAM estates over time (e.g., before/after re-architecture), (ii) set governance thresholds (e.g., maximum acceptable trust density or maximum attack-path length to sensitive assets), (iii) prioritise remediation by focusing on the structural drivers that create reachable high-privilege states, and (iv) support future automation by integrating scoring into CI/CD and cloud security posture workflows [6], [12]. While the proposed model is theoretical at the proposal stage, it is designed to be implementable with common tooling (e.g., graph modelling and Python-based scoring) [6].

2. Literature review

Purpose and scope. This literature review surveys research and practitioner literature on (i) cloud IAM misconfiguration and excessive privilege as breach drivers, (ii) conceptualisations of IAM complexity (e.g., privilege creep, policy sprawl, role explosion), (iii) graph-based

security analysis and attack-path modelling, and (iv) risk assessment frameworks that inform scoring and weighting choices (e.g., NIST SP 800-30 and FAIR). The review is organised thematically to identify key contributions and expose gaps relevant to developing an IAM complexity risk-scoring model [1]–[3], [6], [9], [10], [13].

2.1 Cloud IAM misconfiguration and excessive privilege

Across cloud providers, IAM misconfigurations and excessive permissions are repeatedly highlighted as high-impact security issues because they enable privilege escalation, lateral movement, and direct access to sensitive resources [7], [10], [15]. Practitioner analyses commonly point to overly permissive roles/policies, weak identity hygiene, and unmanaged trust relationships as recurring patterns behind cloud incidents [10]. The persistence of these patterns is often attributed to the scale and pace of cloud change, where temporary exceptions become permanent and where machine identities proliferate alongside human users [9], [10].

Best-practice guidance emphasises least privilege and systematic policy governance, but also acknowledges that achieving least privilege “at scale” is operationally challenging as organisations grow [3]. Recommendations include continuous auditing and tooling-assisted policy generation/validation, including provider-supported analysis of unused permissions and right-sizing opportunities [4], [5]. These discussions motivate the need for measurable indicators of IAM complexity: if complexity rises faster than governance capacity, risk may increase even when organisations intend to follow least-privilege principles [3], [10].

2.2 IAM complexity drivers

IAM complexity can be discussed through several recurring constructs. *Privilege creep* describes gradual accumulation of entitlements beyond current job or workload needs. *Policy sprawl* reflects rapid growth and fragmentation of policy artefacts, often leading to duplicated or inconsistent statements and difficulty predicting effective access. *Role explosion* occurs when organisations create many narrowly scoped roles to follow least privilege, but without consolidation and lifecycle controls those roles can proliferate, overlap, and become opaque [3], [9], [13].

In cloud environments, complexity is amplified by trust relationships and delegation mechanisms (e.g., role assumption and cross-account access), where effective permissions depend on combinations of identity policies, resource policies, and trust policies [6], [12]. Least-privilege deviation can be seen as the distance between granted privileges and the minimal privilege set required for tasks; however, defining “required” privileges is itself a non-trivial problem, and provider tooling therefore often focuses on identifying unused permissions as a practical proxy [4], [5]. A key gap in the reviewed literature is the absence of a unified, IAM-specific set of structural metrics that can be combined into a defensible complexity score suitable for comparison across environments and over time [6], [9], [13].

2.3 Graph-based analysis, trust relationships, and attack paths

Graph-based modelling is widely used in security analysis to represent systems as interconnected entities and relationships, enabling discovery of transitive dependencies and

multi-step exploitation paths. In cloud security, graph paradigms underpin “attack path analysis”, where misconfigurations, permissions, identities, and resources are connected into plausible routes to high-value assets [6], [12]. This approach is directly relevant to IAM, where trust edges and policy edges can combine to form privilege escalation chains that are difficult to detect through list-based review [6], [7].

Prior work and tooling demonstrate feasibility of representing cross-account role trust and “assume role” relationships as graphs and then querying for dangerous paths [6], [12]. Such work motivates candidate metrics for an IAM complexity score, including trust-relationship density, number of alternative paths to privileged roles, shortest-path distance from exposed principals to sensitive assets, and overall attack-path complexity [6]. The literature nonetheless tends to focus on *finding* paths rather than producing a consolidated complexity/risk score that can be tracked and governed [6], [9].

2.4 Risk assessment frameworks

NIST SP 800-30 provides general guidance for conducting risk assessments, framing risk in terms of threat events, vulnerabilities, likelihood, and impact, and recommending structured analysis to prioritise risks [1]. This proposal uses NIST’s framing to justify that IAM complexity can be treated as a contributor to likelihood (more exploitable pathways, higher chance of misconfiguration) and potentially impact (larger blast radius through broad entitlements and dense trust) [1], [10].

FAIR (Factor Analysis of Information Risk) is a widely cited quantitative model that decomposes risk into loss event frequency and loss magnitude and provides a structured taxonomy of contributing factors [2]. FAIR’s value for this proposal is conceptual: it supports the argument that risk quantification should be decomposable and defensible, which aligns with building an IAM complexity score out of explicit factors and weights [2]. However, FAIR does not specify how to derive inputs from IAM structural properties. Therefore, an opportunity exists to connect IAM graph-derived complexity indicators to risk-scoring constructs in a transparent way [2], [6].

2.5 Literature summary and gap statement

In summary, the literature establishes that (i) IAM misconfigurations and excessive privileges are key contributors to cloud incidents [7], [10], [15], (ii) IAM complexity grows through identifiable mechanisms (privilege creep, policy sprawl, role explosion, and trust relationship proliferation) [3], [9], [13], (iii) graph-based approaches are effective for modelling reachability and attack paths [6], [12], and (iv) mature risk frameworks support structured reasoning but are not IAM-specific [1], [2]. The primary gap is the lack of a unified theoretical model that translates IAM structural complexity into a repeatable, measurable risk score [6], [9]. This study positions itself to bridge that gap by defining explicit IAM complexity factors, operationalising them with graph-based representations, and combining them through a weighted scoring approach designed for governance and longitudinal comparison [1], [2], [6].

3. Research Design and Methodology

3.1 Research design

This research adopts a **theoretical/conceptual modelling** design, with elements of **design science** in the sense that it produces an artefact (a scoring model) and evaluates it through conceptual validation. The design is appropriate because the objective at proposal stage is not to measure risk empirically across production cloud environments, but to define a coherent model that (i) specifies construct definitions (complexity factors), (ii) provides a representation (graph model), and (iii) yields a calculable score (weighted aggregation). Conceptual validation through hypothetical scenarios enables early assessment of model behaviour before later empirical testing in an MSc project [1], [2].

3.2 Research methods and materials

3.2.1 Methods

Methodological approach: The study uses a predominantly **quantitative** approach in the sense that it defines numerical indicators (metrics) and combines them into a score. At the same time, the research involves **qualitative** conceptual work to define constructs, justify factor selection, and argue how structural indicators relate to risk concepts. This combination is appropriate because IAM complexity must be operationalised (measured) but also theoretically grounded (explained) within established risk assessment and quantification perspectives [1], [2].

Research strategy: The research will follow a model-building strategy consisting of: (1) factor definition and operationalisation, (2) graph representation design, (3) scoring formulation and weighting, and (4) conceptual validation using hypothetical scenarios [6], [12].

1. **Factor definition and operationalisation:** Define each complexity factor and propose measurable proxies. For example: privilege creep (growth of permissions per identity over time, or breadth of actions/resources granted); policy sprawl (count of policies, average policy size/overlap); role explosion (count of roles per business unit/service, redundancy/overlap); trust relationship density (graph density among principals/roles and trust edges); least-privilege deviation (unused permissions indicators or distance from minimal policy) [4], [5]; attack path complexity (number/length of privilege escalation paths to sensitive resources) [6], [12].
2. **Graph representation:** Model IAM as a directed labelled graph: nodes may include principals (users, service accounts, roles), policy objects, and resources; edges capture relationships such as “has policy”, “allows action”, “assumes role”, “trusts”, and “can access resource”. The exact schema will be selected to support path queries and metric computation, consistent with graph-based cloud attack path analysis approaches [6], [12].
3. **Weighted scoring:** Normalise each factor into a comparable scale (e.g., 0–1 or 0–100), define weightings, and aggregate into an overall score. Weights will be justified using literature and risk logic (factors that directly create privilege escalation paths

may be weighted higher) and aligned to established risk assessment/quantification principles [1], [2].

4. **Data analysis and evaluation:** Evaluate model behaviour via scenario-based tests: monotonicity checks (score increases when risk-relevant complexity increases), sensitivity analysis (how changes in one factor alter total score), and plausibility checks (expert judgement or supervisor feedback on whether the score aligns with perceived risk), including checks against graph-derived attack-path intuition [6].

3.2.2 Materials

Research data: The initial validation will use **hypothetical IAM scenarios** (synthetic configurations) designed to reflect realistic patterns such as cross-role trust chains, overlapping policies, and permission growth over time. Hypothetical scenarios are appropriate at this stage because the project's focus is theoretical model development, and because real IAM datasets may be sensitive or unavailable without formal approvals; additionally, incident and threat reporting highlights that misconfigurations and excessive permissions can be operationally sensitive when tied to real environments [10].

Data collection methods and tools: Synthetic IAM graphs will be created manually and/or programmatically (e.g., parameterised generators that can vary role counts, trust density, and permission breadth). If later extended to empirical validation, data could be collected via cloud provider APIs/exports (e.g., IAM policy documents, role trust policies, and access relationships) subject to organisational permission and ethical review; public datasets may also support controlled experimentation [11].

Ethical considerations: The proposal stage avoids processing real organisational IAM configurations. If real data is used later, it may contain sensitive security information (e.g., privileged roles, resource identifiers, trust policies) and would require appropriate approvals, anonymisation/pseudonymisation of identifiers, secure storage, and minimisation of data retained. The study will follow institutional ethics guidance and the principle of least data exposure; where possible, testing can be performed using synthetic scenarios or publicly available datasets rather than production configurations [11].

Software: Python for metric computation and scoring; graph libraries (e.g., NetworkX) and/or a graph database/query tool for representing IAM relationships; a reference manager (Zotero or Mendeley) for citations; standard word processing tools. Graph-based modelling choices will be informed by prior work applying graphs to cloud tenant attack paths [6] and practitioner tooling that maps IAM trust paths [12].

Hardware: Modelling and scoring can be performed on a standard personal computer for hypothetical scenarios. If later extended to large-scale empirical graphs, additional compute (cloud virtual machines) may be needed to store and query graphs efficiently; GPU resources are not expected to be required.

4. Expected outcome

Anticipated results: The primary expected outcome is a clearly specified theoretical IAM complexity risk-scoring model, including (i) a factor set with definitions and measurable proxies, (ii) a graph representation schema for IAM relationships, and (iii) a scoring function with normalisation and weights. Conceptual validation is expected to show that the score increases consistently as scenario complexity increases (more trust chains, broader policies, more alternative privilege escalation paths).

Implications. The model could provide a foundation for future research that empirically tests whether measured IAM complexity correlates with incidents, audit findings, or misconfiguration rates. Practically, it may support governance by enabling IAM teams to track complexity as a leading indicator and prioritise simplification efforts that reduce reachable high-privilege states.

Limitations and risks: At proposal stage the model is theoretical and validated only using hypothetical scenarios, which may not capture all nuances of real cloud estates. Weight choices may be subjective and require refinement through expert elicitation or empirical calibration. Cloud IAM capabilities evolve rapidly, so factor definitions and metrics may need periodic updates to remain relevant.

5. Timeline

Workstream / milestone	May	Jun	Jul	Aug	Sep	Oct
Literature review (core themes + gap statement)	Start	Complete				
Define IAM complexity factors + candidate metrics		Start	Complete			
Design graph representation (schema) + scenario generator			Start	Complete		
Develop weighted scoring model + normalisation approach			Start	Complete		
Conceptual validation (hypothetical scenarios + sensitivity tests)				Start	Complete	
Write-up, formatting, and supervisor/peer feedback					Start	Complete

7. References

IEEE referencing note. Cite sources in-text using IEEE numeric citations in square brackets (e.g., [1], [2]–[4]) and list full references below in the order they first appear in the text.

- [1] National Institute of Standards and Technology (NIST), *Guide for Conducting Risk Assessments*, NIST Special Publication 800-30 Rev. 1, Gaithersburg, MD, USA, 2012.
- [2] The FAIR Institute, *Factor Analysis of Information Risk (FAIR) Standard*, Version 3.0, 2025.
- [3] Amazon Web Services (AWS), “IAM best practices,” 2024. Accessed: Apr. 20, 2026.
- [4] S. Mbappe and M. Ramesh, “Refine unused access using IAM Access Analyzer recommendations,” *AWS Security Blog*, Sep. 18, 2024. Accessed: Apr. 20, 2026.
- [5] L. D’Antoni and C. Sung, “New AWS tool recommends removal of unused permissions,” *Amazon Science*, Dec. 19, 2024. Accessed: Apr. 30, 2026.
- [6] M. Elmiger, M. Lemoudden, N. Pitropakis, and W. J. Buchanan, “Start thinking in graphs: Using graphs to address critical attack paths in a Microsoft cloud tenant,” *International Journal of Information Security*, vol. 23, pp. 467–485, 2024, doi: 10.1007/s10207-023-00751-6.
- [7] I. B. Haimed, M. Albahar, and A. Alzubaidi, “Exploiting misconfiguration vulnerabilities in Microsoft’s Azure Active Directory for privilege escalation attacks,” *Future Internet*, vol. 15, no. 7, Art. no. 226, 2023, doi: 10.3390/fi15070226.
- [8] A. Gougolidis, A. Kagia, and V. C. Hu, “Model checking access control policies: A case study using Google Cloud IAM,” arXiv:2303.16688, 2023.
- [9] J. Glöckler, J. Sedlmeir, M. Frank, and G. Fridgen, “A systematic review of identity and access management requirements in enterprises and potential contributions of self-sovereign identity,” *Business & Information Systems Engineering*, vol. 66, pp. 421–440, 2024, doi: 10.1007/s12599-023-00830-x.
- [10] Palo Alto Networks Unit 42, *Cloud Threat Report, Volume 6*, Apr. 12, 2022. Accessed: Apr. 21, 2026.
- [11] A. M., V. S. Salanke, and S. Nagasundari, “Misconfiguration detection dataset for cloud IAM and APIs,” *IEEE DataPort*, 2025, doi: 10.21227/m0h8-en82.
- [12] A. Kallio, “IAMGraph: Mapping cross-account attack paths in AWS environments,” *WithSecure Labs*, 2026. Accessed: Apr. 21, 2026.
- [13] D. A. Shafiq *et al.*, “Access control models in cloud computing,” *IEEE Access*, vol. 8, pp. 219456–219475, 2020.
- [14] M. Ali *et al.*, “Security in cloud computing: Opportunities and challenges,” *Information Sciences*, vol. 305, pp. 357–383, 2015.

[15] D. J. Millward, M. J. Reed, and N. Olaniyi, “CACLE—Automated mitigation for misconfiguration vulnerabilities in cloud systems,” in *Proc. 7th Int. Conf. Cloud and Big Data Computing (ICCBDC)*, Manchester, U.K., Aug. 2023, pp. 1–9, doi: 10.1145/3616131.3616142.