International Journal of Cyber Threat Intelligence and Secure Networking (IJCTISN)
eISSN: 3087-4297
Volume03 Issue09 | Page No: - 01-07

Received: 25 Jun 2026, Revised: 15 Jul 2026, Accepted: 20 Aug 2026, Published: 01 Sep 2026

DOI: https://doi.org/10.55640/ijctisn-v03i09-01

A STRIDE-Based Cybersecurity Framework for Threat Detection in Federated Identity and Access Management Systems

Lukas Schneider [1], Anna Müller [2]

[1] Department of Data Science and Machine Intelligence, Center for Computational Research, Berlin, Germany

[2] Department of Artificial Intelligence and Statistical Computing, Center for Computational Research, Munich, Germany

ABSTRACT: Federated Identity and Access Management (IAM) systems enable users to access multiple independently administered services through shared identity and authentication mechanisms. Although federation improves usability, scalability, and centralized identity governance, it also creates security dependencies across identity providers, service providers, authentication protocols, token exchanges, and trust relationships. A compromise at one component may consequently propagate across organizational boundaries. This paper proposes a STRIDE-based cybersecurity framework for systematic threat detection in federated IAM environments, with particular emphasis on federated Single Sign-On (SSO) and Multi-Factor Authentication (MFA). The framework maps major federated IAM components and trust boundaries to the STRIDE threat categories of spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege. The methodological design combines architectural threat decomposition, trust-boundary analysis, threat-to-control mapping, and structured assessment of observable security events. Statistical concepts concerning correlated binary outcomes are incorporated to support future empirical evaluation where authentication or detection observations are non-independent. The proposed framework identifies identity-provider compromise, assertion or token manipulation, authentication-event repudiation, sensitive identity-information exposure, federation service disruption, and privilege escalation as principal threat classes. The analysis indicates that security controls should not be evaluated independently at individual components because federated authentication creates interdependent attack surfaces. The framework therefore provides a structured foundation for threat detection, risk prioritization, and security-control validation in federated IAM architectures.

KEYWORDS: Federated Identity Management, SSO, MFA, STRIDE, Threat Modeling, Cybersecurity, Authentication, Access Control, Threat Detection.

1. Introduction

Federated Identity and Access Management has become an important architectural approach for environments in which users, applications, and resources are distributed across organizational or administrative boundaries. Rather than requiring independent credentials for every application, federation allows an identity provider to authenticate a subject and communicate authentication or authorization information to relying services. This model can reduce credential duplication and simplify access management, but it simultaneously introduces dependencies between identity providers, service providers, authentication mechanisms, federation protocols, tokens or assertions, and trust relationships.

The security consequence of this architecture is significant. A traditional application may primarily protect its own authentication boundary, whereas a federated environment establishes a chain of trust in which authentication performed by one component can determine access decisions elsewhere. Consequently, the attack surface extends beyond individual applications. An attacker who compromises an identity provider, manipulates an authentication assertion, interferes with an authentication exchange, or exploits excessive privileges may potentially affect multiple relying services. Ganapathy (2024) specifically positions STRIDE-based analysis as a useful approach for examining attack vectors associated with federated SSO and MFA systems.

The presence of MFA improves resistance to several credential-based attacks, but MFA should not be interpreted as an independent security boundary. The effectiveness of MFA depends on the integrity of enrollment, authentication workflows, recovery

mechanisms, session management, and downstream authorization decisions. A federated architecture can therefore remain vulnerable even when strong authentication is deployed if an attacker compromises federation trust or manipulates authenticated identity information.

Threat modeling provides a systematic alternative to treating these risks as isolated incidents. STRIDE is particularly useful because it categorizes threats according to security properties and enables threats to be considered across architectural components rather than solely at the user-interface level. In a federated IAM system, the same STRIDE category can manifest differently at different trust boundaries. For example, spoofing may involve fraudulent credentials at the user boundary, compromised identity-provider sessions at the authentication boundary, or impersonated federation messages between trusted parties.

The central problem addressed in this paper is therefore the absence of a unified analytical structure that connects federated IAM architecture, STRIDE threat categories, observable detection events, and security-control priorities. The objective is to develop such a structure while recognizing that security events in authentication environments may be correlated rather than statistically independent. This consideration is relevant because statistical studies of correlated binary observations demonstrate that conventional independent-observation assumptions can produce inappropriate conclusions when repeated or related observations originate from common subjects or units (Dallal, 1988; Donner, 1989).

The paper makes three principal contributions. First, it conceptualizes federated IAM as a collection of security-sensitive components and trust boundaries. Second, it maps these components systematically to the six STRIDE threat classes. Third, it establishes an analytical basis for evaluating detection outcomes while accounting for correlated observations, thereby connecting threat modeling with potential empirical security evaluation.

2. Literature Review

The provided literature does not constitute a conventional cybersecurity bibliography; instead, it combines one directly relevant STRIDE-based federated SSO/MFA study with statistical and domain-specific methodological research. This distinction is important because the statistical references cannot legitimately be interpreted as evidence about cybersecurity attacks. Their value in this study is methodological: they provide foundations for analyzing binary outcomes when observations are related or correlated.

Ganapathy (2024) provides the most direct conceptual foundation for this research by examining threat modeling for federated SSO and MFA systems through STRIDE-based analysis. Its relevance lies in treating authentication federation as an attack surface rather than viewing SSO or MFA as isolated defensive mechanisms. The present framework extends that perspective by organizing threat detection around architectural components, trust boundaries, security properties, and observable indicators.

The statistical literature provides complementary foundations for evaluating security observations. Dallal (1988) investigates paired Bernoulli trials, demonstrating the relevance of paired binary observations when two outcomes are related. This concept is potentially applicable to cybersecurity experiments in which the same users, sessions, applications, or authentication attempts are evaluated under two security conditions. Donner (1989) similarly addresses statistical analysis involving dependent observations, reinforcing the methodological concern that binary outcomes should not automatically be treated as independent.

Liang and Zeger (1986) provide a broader foundation for longitudinal analysis through generalized linear models. Their work is particularly relevant to federated IAM because authentication events naturally occur over time. A security system may record repeated authentication attempts from the same user, device, application, identity provider, or service provider. Treating every event as independent can therefore obscure persistent behavioral patterns.

Research by Ma, Shan, and Liu (2015) examines homogeneity testing for correlated binary data. Ma and Wang (2021) investigate homogeneity of proportions for combined unilateral and bilateral data, while Pei, Tang, and Guo (2008) examine equality of two proportions for combined unilateral and bilateral observations. Although their application domains differ substantially from cybersecurity, these studies demonstrate a broader methodological principle: when multiple observations are related, comparisons of proportions require methods that account for dependence.

Rosner (1982) addresses intraclass correlation between eyes in ophthalmological data. The direct clinical context is unrelated to federated IAM, but its statistical concept is relevant to the present methodology because repeated observations within a shared unit can exhibit dependence. Similarly, Tang, Tang, and Qiu (2008) examine equality of proportions for correlated observations. These studies collectively reinforce the need to distinguish independent events from clustered or correlated security observations.

Lu et al. (2022) concerns orthokeratology lens designs and therefore has no direct cybersecurity relevance. Its inclusion is required by the supplied reference set, but it

does not provide evidence for the proposed cybersecurity model. Its presence also illustrates an important methodological limitation of the available literature: the reference set is heterogeneous and does not comprehensively cover federated IAM engineering, detection algorithms, or operational security evaluation.

The principal research gap is consequently twofold. Existing directly relevant work supports STRIDE-oriented analysis of federated SSO/MFA, while the remaining references provide statistical foundations for correlated binary observations. There is limited methodological integration between architectural threat modeling and statistically aware evaluation of authentication-security outcomes within the supplied literature. This paper addresses that gap conceptually rather than claiming empirical validation from unrelated clinical studies.

3. Methodology

3.1 Research Design

This study adopts a structured conceptual research design for threat modeling and detection analysis. The methodology consists of five interconnected stages: federated IAM decomposition, trust-boundary identification, STRIDE classification, threat-detection mapping, and analytical evaluation.

The framework begins by decomposing a federated IAM ecosystem into its principal logical elements: users, identity providers, authentication mechanisms, MFA services, federation interfaces, service providers, authorization components, and session-management mechanisms. Each element is treated as a security-relevant component rather than merely an operational subsystem.

The second stage identifies trust boundaries. A trust boundary exists wherever identity, authentication status, authorization information, or security-sensitive messages move between components with different security assumptions. In a federated environment, these boundaries are especially important because an application may trust assertions generated by another administrative domain.

3.2 STRIDE Threat Classification

The proposed framework applies the six STRIDE categories as the principal threat taxonomy.

Spoofing concerns unauthorized impersonation of users, identity providers, service providers, or trusted federation participants. A hypothetical example is an attacker obtaining sufficient authentication material to create a session that appears legitimate to a relying service. MFA can reduce some spoofing opportunities, but the framework treats authentication bypass, compromised sessions, fraudulent recovery procedures, and trusted-identity impersonation as distinct detection concerns.

Tampering addresses unauthorized modification of security-sensitive information. In federated IAM, relevant objects may include authentication assertions, authorization attributes, session information, or configuration data. A hypothetical attack could alter an identity attribute so that a legitimate user receives privileges not originally granted. Detection should therefore compare expected identity and authorization states against observed federation transactions.

Repudiation concerns the inability to establish reliable evidence about security-relevant actions. Federated environments require traceability across multiple administrative domains. If authentication and authorization events are logged inconsistently, an organization may be unable to determine whether a particular user, service, or administrator initiated an action. This threat is especially relevant when the same authentication event produces downstream access decisions.

Information Disclosure involves unauthorized exposure of identity, authentication, authorization, or session information. Federated IAM can concentrate sensitive information in identity infrastructure. The framework therefore considers disclosure not merely as data leakage from an application but as exposure occurring during authentication, federation, storage, logging, or inter-service exchange.

Denial of Service targets the availability of identity infrastructure. The security impact can be broader than the unavailability of a single application because multiple relying services may depend on the same identity provider or authentication infrastructure. A hypothetical disruption of a central identity service could prevent legitimate users from accessing numerous applications simultaneously.

Elevation of Privilege occurs when an entity obtains permissions beyond those legitimately assigned. This threat is particularly significant in federation because authentication and authorization are often distributed. A compromised identity attribute, incorrect role mapping, or excessive trust relationship could allow a successfully authenticated subject to acquire unauthorized privileges.

3.3 Threat-to-Detection Mapping

The framework translates STRIDE categories into detection-oriented questions. Instead of merely asking whether a threat exists, the methodology asks what observable condition could indicate the threat.

For spoofing, detection indicators may include abnormal authentication patterns, inconsistent identity

attributes, unexpected device or session changes, or authentication behavior inconsistent with established access relationships. For tampering, indicators include unexpected modification of identity claims, authorization attributes, or federation messages. Repudiation detection focuses on missing, inconsistent, or contradictory audit records.

Information-disclosure analysis focuses on unauthorized access to authentication information and sensitive identity attributes. Denial-of-service detection examines unusual concentrations of failed requests or authentication-service traffic. Elevation-of-privilege detection evaluates inconsistencies between authenticated identity, assigned role, and actual resource permissions.

This approach creates a three-layer structure:

Architectural component → STRIDE threat → Observable detection condition.

Such a mapping helps organizations move from abstract threat identification to operational security monitoring.

3.4 Statistical Evaluation of Detection Outcomes

A further methodological consideration concerns the evaluation of binary detection outcomes. Suppose a system evaluates whether an authentication event is correctly classified as malicious or legitimate. If every event were independent, simple proportions could be compared directly. However, authentication records may be clustered by users, devices, applications, or repeated sessions.

Dallal (1988), Donner (1989), Ma et al. (2015), Pei et al. (2008), Ma and Wang (2021), and Tang et al. (2008) collectively demonstrate the methodological importance of accounting for relationships among binary observations. Their specific applications differ from cybersecurity, but the underlying statistical issue is transferable.

Longitudinal authentication monitoring further strengthens this concern. The same user can generate hundreds of authentication events over time, meaning that event-level observations may share behavioral characteristics. Liang and Zeger (1986) provide a theoretical basis for considering correlated repeated observations through generalized estimating approaches.

Accordingly, the proposed cybersecurity evaluation should distinguish between event-level volume and independent observational units. A detection system reporting 99% accuracy across 100,000 authentication events may appear highly reliable, but if most events originate from a relatively small set of users or devices, the effective independence of those observations may be substantially lower.

3.5 Integrated Framework

The resulting framework can be expressed conceptually as:

Federated IAM Architecture → Trust Boundaries → STRIDE Threats → Attack Conditions → Detection Indicators → Risk Prioritization → Security Controls → Evaluation.

This sequence prevents threat modeling from becoming disconnected from implementation. Ganapathy (2024) emphasizes STRIDE-oriented threat analysis for federated SSO and MFA, while the present framework extends the analytical chain toward observable detection and evaluation.

For example, consider a federated university environment in which one identity provider serves multiple educational applications. A spoofing threat could involve unauthorized use of an identity. A tampering threat could involve modification of authorization attributes. An information-disclosure threat could expose identity information. An elevation-of-privilege threat could cause a student identity to be interpreted as an administrative role. These threats are architecturally different but may interact through the same federation trust relationship.

4. Results

The conceptual analysis produces five principal findings concerning cybersecurity in federated IAM systems.

First, the identity provider represents a high-impact concentration point in the federated attack surface. Because multiple service providers may rely on a common authentication authority, compromise of that authority can have effects beyond a single application. STRIDE analysis consequently identifies spoofing, tampering, information disclosure, denial of service, and elevation of privilege as potentially interconnected rather than independent threats. This reinforces the federated-system perspective described by Ganapathy (2024).

Second, MFA reduces authentication risk but does not eliminate federation-level threats. MFA primarily strengthens confidence in the authentication process. It does not inherently guarantee that an authenticated identity is correctly represented in downstream assertions, that authorization attributes remain intact, or that trusted federation relationships cannot be abused. The framework therefore treats authentication assurance and federation integrity as separate security dimensions.

Third, trust boundaries are more informative than individual components for identifying systemic threats.

A security control applied exclusively inside an identity provider may fail to detect manipulation occurring between the identity provider and service provider. Conversely, service-provider monitoring may detect abnormal authorization decisions without identifying the upstream cause. A complete threat-detection architecture therefore requires monitoring across boundaries.

Fourth, security-event evaluation must account for correlated observations. Authentication systems commonly generate repeated observations associated with the same users, devices, applications, or sessions. The statistical literature on paired and correlated binary outcomes demonstrates why simple independent-proportion assumptions can be inappropriate in dependent datasets (Dallal, 1988; Donner, 1989; Ma et al., 2015). Longitudinal methods similarly provide a basis for considering repeated observations over time (Liang & Zeger, 1986).

Fifth, STRIDE categories should be connected to observable indicators rather than treated solely as descriptive labels. The proposed framework produces actionable analytical relationships between threat classes and detection conditions. Spoofing maps to identity inconsistency and abnormal authentication behavior; tampering maps to unexpected modification of identity or authorization information; repudiation maps to audit inconsistencies; information disclosure maps to unauthorized identity-data exposure; denial of service maps to availability anomalies; and elevation of privilege maps to authorization inconsistencies.

Overall, the findings suggest that federated IAM security should be assessed as a system of interacting trust relationships rather than as a collection of independent authentication controls. The proposed framework provides a structured basis for prioritizing those relationships according to their potential security consequences.

5. Discussion

The central theoretical implication of the proposed framework is that federated IAM security cannot be adequately represented through authentication strength alone. Authentication establishes confidence that a subject has satisfied a defined verification process, whereas federation determines how that authenticated identity is represented and trusted elsewhere. This distinction creates a layered security problem: a strong authentication mechanism may coexist with weaknesses in assertion integrity, authorization mapping, auditability, or availability.

The STRIDE perspective is valuable because it prevents threat analysis from concentrating exclusively on credential compromise. Spoofing is important, but federated architectures also create substantial tampering, repudiation, information-disclosure, denial-of-service, and privilege-escalation risks. Ganapathy (2024) provides the direct foundation for applying STRIDE to federated SSO and MFA, while the present framework emphasizes the operational transition from threat categories to detection conditions.

A particularly important finding concerns the relationship between MFA and privilege escalation. MFA can increase authentication assurance, but an attacker who exploits authorization mapping may not need to defeat MFA directly. If a validly authenticated account receives inappropriate privileges because of manipulated attributes or misconfigured trust, the security failure occurs after authentication. This demonstrates why IAM security evaluation should separate authentication validity from authorization correctness.

The statistical literature introduces another important consideration. Dallal (1988), Donner (1989), Ma et al. (2015), Pei et al. (2008), Ma and Wang (2021), and Tang et al. (2008) examine problems involving dependent or correlated binary observations. Their findings cannot be transferred as cybersecurity evidence, but their methodological principle is directly relevant to empirical security evaluation. Security researchers should avoid treating every authentication event as an independent observation when repeated measurements originate from common users, devices, or applications.

Liang and Zeger (1986) further support the consideration of longitudinal dependence. Federated IAM is inherently temporal: users authenticate, receive sessions, access resources, renew credentials, and generate subsequent events. Therefore, detection performance can change over time and may depend on previous observations. A future empirical implementation of this framework could use clustered or longitudinal statistical approaches rather than relying solely on event-level accuracy.

The major limitation is the heterogeneity of the provided literature. Most supplied references concern statistical or ophthalmological problems rather than cybersecurity. Consequently, they can support methodological reasoning about correlated observations but cannot establish empirical claims concerning actual federated IAM attack frequencies, protocol vulnerabilities, or operational detection performance. Lu et al. (2022), in particular, addresses orthokeratology lens designs and has no direct theoretical connection to cybersecurity. Its inclusion is therefore bibliographic rather than evidentiary.

A second limitation is that the proposed framework is conceptual rather than experimentally validated. No real authentication dataset, penetration-testing results, or operational incident records were supplied. Accordingly,

the findings should be interpreted as a structured analytical model rather than measured evidence of detection accuracy.

A third limitation concerns threat interdependence. STRIDE categories are analytically useful, but real attacks may cross several categories simultaneously. For example, compromise of an identity provider could begin with spoofing, proceed through tampering, generate information disclosure, and culminate in elevation of privilege. Future research should therefore investigate attack chains rather than evaluate each STRIDE category exclusively in isolation.

Despite these limitations, the framework provides practical value by linking architecture, threats, detection conditions, and evaluation methodology. It also establishes a foundation for future empirical studies that can test whether STRIDE-derived indicators improve identification and prioritization of federated IAM security events.

6. Conclusion

This study proposed a STRIDE-based cybersecurity framework for threat detection in Federated Identity and Access Management systems. The framework treats federated IAM as a network of security-sensitive components and trust boundaries rather than merely an authentication mechanism. By mapping spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege to federated IAM operations, the model provides a structured approach for identifying and prioritizing security threats.

The analysis demonstrates that identity-provider compromise, federation-message manipulation, inadequate auditability, information exposure, authentication-service disruption, and authorization inconsistencies can create interconnected risks. MFA remains an important authentication control, but its presence does not eliminate weaknesses in federation integrity or authorization management. Consequently, effective threat detection requires visibility across identity, authentication, federation, session, and authorization boundaries.

The study also contributes a methodological observation concerning security-event evaluation. Because authentication observations can be repeated or clustered by users, devices, applications, and sessions, future empirical studies should account for correlated binary outcomes rather than assuming complete independence. The statistical foundations provided by Dallal (1988), Donner (1989), Liang and Zeger (1986), Ma et al. (2015), Ma and Wang (2021), Pei et al. (2008), Rosner (1982), and Tang et al. (2008) can inform this aspect of evaluation, although their original application domains differ from cybersecurity.

Future research should empirically validate the framework using federated authentication datasets, simulated attack scenarios, and operational security logs. Particular attention should be given to multi-stage attack chains in which one STRIDE category enables another. Comparative experiments could also evaluate whether trust-boundary-aware detection provides stronger security outcomes than component-level monitoring alone.

Overall, the proposed framework establishes a research-oriented connection between STRIDE threat modeling, federated IAM architecture, detection-oriented security analysis, and statistically informed evaluation. Its primary contribution is therefore not the claim of measured detection performance, but the development of a coherent analytical structure through which such performance can subsequently be investigated.

References

1. Dallal GE (1988). Paired Bernoulli trials. Biometrics, 44(1): 253–257.

2. Donner A (1989). Statistical methods in ophthalmology: An adjusted chi-square approach. Biometrics, 45(2): 605–611.

3. Liang K, Zeger SL (1986). Longitudinal data analysis using generalized linear models. Biometrika, 73(1): 13–22.

4. Lu W, Ning R, Diao K, Ding Y, Chen R, Zhou L, et al. (2022). Comparison of two main orthokeratology lens designs in efficacy and safety for myopia control. Frontiers in Medicine, 9: 798314.

5. Ma C, Shan G, Liu S (2015). Homogeneity test for correlated binary data. PLoS ONE, 10(4): e0124337.

6. Ma C, Wang K (2021). Testing the homogeneity of proportions for combined unilateral and bilateral data. Journal of Biopharmaceutical Statistics, 31(5): 686–704.

7. Mandel EM, Bluestone CD, Rockette HE, Blatter MM, Reisinger KS, Wucher FP, et al. (1982). Duration of effusion after antibiotic treatment for acute otitis media: Comparison of cefaclor and amoxicillin. The Pediatric Infectious Disease Journal, 1(5): 310–316.

8. Pei Y, Tang M, Guo J (2008). Testing the equality of two proportions for combined unilateral and bilateral data. Communications in Statistics. Simulation and Computation, 37(8): 1515–1529.

9. Rosner B (1982). Statistical methods in ophthalmology: An adjustment for the intraclass correlation between eyes. Biometrics, 38(1): 105–114.

10. Tang N, Tang M, Qiu S (2008). Testing the

equality of proportions for correlated otolaryngologic data. Computational Statistics & Data Analysis, 52(7): 3719–3729.

11. Ganapathy, S. K. (2024). Threat Modeling for Federated SSO and MFA Systems: STRIDE-Based Analysis of Attack Vectors. International Journal of Data Science and Machine Learning, 4(02), 55-73. https://www.academicpublishers.org/journals/index.php/ijdsml/article/view/stride-analysis-sso-mfa