

Risk-Based Security Assessment of Federated Authentication Systems: STRIDE Analysis of Identity, Session, and Access Threats

Nuwan Perera

Department of Data Science and Intelligent Computing, Institute of Advanced Computational Studies, Colombo, Sri Lanka

Ishara Fernando

Department of Machine Learning and Data Analytics, Institute of Advanced Computational Studies, Kandy, Sri Lanka

ABSTRACT

Federated authentication systems have become an important architectural mechanism for enabling users to access multiple services through shared identity infrastructure. Their security, however, depends on the integrity of identity assertions, authentication exchanges, session-management mechanisms, and authorization decisions distributed across multiple trust domains. This paper presents a risk-based security assessment of federated authentication systems using the STRIDE threat-modeling framework, with particular emphasis on identity, session, and access-related threats. The methodology combines system decomposition, trust-boundary analysis, STRIDE categorization, risk prioritization, and analytical interpretation of attack surfaces. A conceptual risk model is developed in which threats are assessed according to their potential effect on identity integrity, confidentiality, availability, authentication assurance, and authorization correctness. The study further considers how machine-learning and statistical modeling concepts can support systematic threat prioritization and anomaly-oriented analysis. Existing methodological foundations concerning deep learning, statistical modeling, nonlinear optimization, and machine-learning estimation provide theoretical support for the analytical treatment of complex security data (Alzubaidi et al., 2021; Coleman and Li, 1994; Farrell et al., 2021). The assessment indicates that identity spoofing, authentication-message tampering, session compromise, privilege escalation, and denial-of-service conditions represent particularly significant risk classes in federated environments. The analysis demonstrates that STRIDE is most effective when combined with risk-based prioritization rather than being used as a simple checklist. The resulting framework provides a structured basis for identifying critical trust-boundary failures and prioritizing defensive controls across federated authentication architectures.

KEYWORDS: Federated Authentication, Single Sign-On, Multi-Factor Authentication, STRIDE, Threat Modeling, Identity Security, Session Security, Access Control, Risk Assessment

INTRODUCTION

Federated authentication changes the traditional security model by allowing an identity established within one security domain to be trusted by multiple relying services. Instead of independently authenticating users at every application, a federated architecture typically distributes authentication responsibilities among identity providers, authentication services, relying applications, and supporting infrastructure. This architectural approach can reduce authentication duplication and improve usability, but it also concentrates security dependencies around identity assertions and trust relationships.

The security implications become more significant when authentication information crosses organizational or technological boundaries. A compromise of an identity provider, authentication transaction, session mechanism, or authorization mapping can potentially affect multiple relying services simultaneously. Consequently, security assessment must examine not only individual components but also the interactions and trust assumptions connecting them.

Threat modeling provides a systematic means of examining these interactions. STRIDE is particularly useful because it categorizes threats according to spoofing, tampering,

repudiation, information disclosure, denial of service, and elevation of privilege. In federated authentication, these categories can be mapped to concrete security concerns involving identity impersonation, modification of authentication information, manipulation of session state, disclosure of credentials or tokens, disruption of authentication services, and unauthorized privilege acquisition. Ganapathy (2024) demonstrates the relevance of STRIDE-based analysis specifically to federated SSO and MFA systems, establishing a direct conceptual foundation for examining attack vectors across authentication components.

The research problem addressed in this paper is therefore the difficulty of translating the distributed attack surface of federated authentication into a structured and prioritized security assessment. A conventional vulnerability inventory may identify individual weaknesses without adequately representing their position within trust relationships. A risk-based STRIDE approach instead evaluates how a threat affects the security properties of identity, session, and access-management processes.

The primary objective is to develop a structured analytical model for assessing security threats in federated authentication systems. The paper specifically investigates three closely connected domains: identity threats, session threats, and access-control threats. A second objective is to demonstrate how risk prioritization can improve the practical value of STRIDE by distinguishing high-impact threats from lower-impact findings. A third objective is to examine how analytical and machine-learning concepts from the supplied literature can provide methodological support for handling complex security observations.

The scope is intentionally focused on conceptual security assessment rather than implementation of a particular federation protocol. The proposed analysis is applicable to architectures containing identity providers, authentication services, relying applications, authentication factors, session-management components, and authorization mechanisms.

LITERATURE REVIEW

The supplied literature provides several complementary theoretical foundations for developing a risk-based analytical approach, although the references are predominantly concerned with machine learning, statistics, optimization, and neural-network modeling rather than federated authentication itself.

Alzubaidi et al. (2021) provide a broad review of deep-learning concepts, architectures, challenges, and applications. Their work is relevant to this research because modern security analytics increasingly involves complex

and high-dimensional observations. A federated authentication environment can generate heterogeneous security events involving authentication failures, token transactions, session changes, access requests, and privilege modifications. Although deep learning itself is not required for STRIDE analysis, the conceptual ability of learning systems to process complex patterns provides a potential analytical extension for threat prioritization.

Farrell, Liang, and Misra (2021) examine deep neural networks for estimation and inference. Their contribution is relevant to the distinction between observable security events and inferred security risk. In federated authentication, an individual event may appear benign while its relationship with other events indicates a broader attack pattern. A risk-based methodology therefore benefits from distinguishing raw observations from analytical conclusions.

The statistical foundations represented by Fan (2018) and Friedman (1991) provide additional support for modeling nonlinear relationships. Security risk is rarely a simple linear function of the number of authentication failures or access attempts. For example, a sequence involving an unusual authentication event followed by session manipulation and privilege escalation may produce a substantially greater risk than the sum of the individual events considered independently. Nonlinear modeling concepts therefore offer theoretical justification for treating threat severity as context dependent.

Friedman, Tibshirani, and Hastie (2001) provide a broader statistical-learning foundation for prediction, inference, and data-driven modeling. Their work supports the conceptual separation between features, observations, predictive interpretation, and model-based decision making. Applied to security assessment, this suggests that threat prioritization can incorporate multiple characteristics rather than relying on a single severity indicator.

Goodfellow, Bengio, and Courville (2016) provide a foundational treatment of deep learning and representation learning. The relevance to this research lies primarily in the possibility of representing complex security states through structured features. A federated authentication transaction can be represented through attributes such as authentication method, identity context, session state, requesting service, privilege level, and temporal behavior. Such representations can support future automated security analytics, although this paper remains focused on STRIDE-based assessment rather than constructing a deep-learning model.

Coleman and Li (1994) address convergence of reflective Newton methods for large-scale nonlinear minimization subject to bounds. Although the work is not a cybersecurity

study, its optimization perspective is relevant to risk prioritization problems in which multiple constrained variables must be considered. Security decisions similarly operate under constraints involving available controls, operational costs, service availability, and acceptable risk.

Muir (2024), as cited in the supplied material, describes Adam stochastic-gradient-descent optimization. This provides another methodological connection between optimization and security analytics, particularly where automated risk scoring or machine-learning-based detection models may require iterative parameter optimization.

Finally, Ganapathy (2024) provides the most directly aligned conceptual reference, addressing threat modeling for federated SSO and MFA systems using STRIDE-based analysis of attack vectors. The present paper extends that perspective by emphasizing risk prioritization across identity, session, and access domains rather than treating STRIDE categories as independent checklist items.

The principal research gap is therefore apparent: the supplied literature provides strong analytical foundations for modeling, optimization, and threat categorization, but there is limited integration of these concepts into a unified risk-based assessment specifically organized around the lifecycle of federated identity. This paper addresses that gap conceptually by combining system decomposition, STRIDE classification, and risk-oriented prioritization.

METHODOLOGY

3.1 Research Design

This study adopts a conceptual, framework-driven research methodology. The objective is not to measure vulnerabilities in a particular deployed federation but to construct a repeatable procedure for identifying and prioritizing threats within federated authentication architectures.

The methodology consists of five sequential stages: system decomposition, trust-boundary identification, STRIDE threat mapping, risk scoring, and analytical interpretation.

System decomposition identifies the principal functional elements participating in authentication and access decisions. These include the user, identity provider, authentication mechanism, relying application, session-management layer, authorization mechanism, and communication channels connecting them.

Trust-boundary identification then determines where security assumptions change. For example, a trust boundary exists when an authentication assertion moves from an identity provider to a relying application. Another may exist between a user device and the authentication service. These

boundaries are important because threats frequently emerge where one component accepts information generated by another component.

STRIDE classification is subsequently applied to each relevant component and interaction. The analysis focuses particularly on three dimensions: identity, session, and access.

3.2 Identity Threat Analysis

Identity threats concern the authenticity and integrity of the principal requesting access. Spoofing is the dominant STRIDE category in this domain because an attacker may attempt to impersonate a legitimate user, identity provider, or authentication component.

A hypothetical example involves an attacker acquiring authentication material and presenting it as evidence of a legitimate identity. The security failure is not merely credential compromise; it is the incorrect establishment of trust between the relying service and an unauthorized actor.

Tampering threats are also significant. If identity assertions or authentication messages can be altered, the relying application may make an authorization decision based on incorrect identity attributes. A modified identity attribute could, for example, transform an ordinary user context into an administrative context.

The assessment therefore evaluates whether identity information remains trustworthy across its complete transmission and processing lifecycle. Ganapathy (2024) similarly positions federated SSO and MFA as systems requiring analysis of multiple attack vectors rather than isolated credential protection.

3.3 Session Threat Analysis

Authentication does not terminate when the user successfully proves identity. A subsequent session represents an ongoing security relationship between the user and application.

Session threats therefore include unauthorized reuse, manipulation, disclosure, or persistence of session information. Within STRIDE, spoofing and tampering are particularly relevant, while information disclosure can expose session-related secrets that facilitate unauthorized access.

A hypothetical example is an attacker obtaining a valid session artifact and using it to access an application without repeating the original authentication process. The resulting threat demonstrates an important distinction: strong initial authentication does not automatically guarantee strong session security.

Session risk should consequently be evaluated using factors such as session lifetime, state transitions, reauthentication requirements, privilege changes, and termination behavior.

3.4 Access and Privilege Threat Analysis

Access threats concern whether authenticated identities receive exactly the privileges they are authorized to possess. Elevation of privilege is therefore the central STRIDE category.

An example is a federated user whose authenticated identity is correctly established but whose attributes are incorrectly mapped to an administrative role. In such a case, authentication succeeds while authorization fails.

This distinction is fundamental to federated security assessment. Authentication establishes who the user is; authorization determines what that user may do. A secure federation must therefore preserve both identity integrity and privilege integrity.

3.5 Risk Prioritization Model

A risk-based assessment requires more than threat identification. Each identified threat must be prioritized according to its potential impact and likelihood.

A conceptual risk score can be represented as:

$$\text{Risk} = \text{Likelihood} \times \text{Impact}$$

Likelihood represents the estimated plausibility of successful exploitation, while impact represents the consequences if exploitation succeeds. Impact can be evaluated across identity compromise, confidentiality loss, integrity degradation, service disruption, and unauthorized privilege acquisition.

A more detailed analytical representation can assign weighted dimensions:

$$R = w_1I + w_2S + w_3A + w_4C + w_5D$$

where I represents identity impact, S session impact, A access-control impact, C confidentiality impact, and D availability impact. The weights reflect organizational priorities.

The statistical-learning literature suggests why such multidimensional assessment can be useful. Complex relationships between variables may not be adequately represented by simplistic linear assumptions (Fan, 2018; Friedman, 1991). Similarly, optimization principles can support systematic selection of risk priorities under bounded operational constraints (Coleman and Li, 1994).

3.6 Analytical Procedure

The final procedure evaluates every major federation component against applicable STRIDE categories. Each threat is then assigned an impact level and likelihood level. Threats affecting identity establishment, session continuity, or privilege boundaries receive additional analytical attention because compromise in these areas can propagate across multiple relying services.

The methodology is intentionally technology-neutral. This improves general applicability but also creates a limitation: specific protocol behaviors and implementation weaknesses cannot be established without deployment-specific evidence.

RESULTS

The conceptual assessment identifies three dominant security zones in federated authentication: identity establishment, session continuity, and access authorization. These zones are strongly interconnected, meaning that compromise in one can amplify risks in another.

The identity layer demonstrates the highest systemic significance because it establishes the foundation upon which subsequent authorization decisions depend. Spoofing threats are consequently particularly important. An attacker who successfully impersonates a trusted identity can potentially bypass multiple downstream security controls. Tampering with identity information represents a second major concern because the relying application may accept manipulated attributes as legitimate. The analysis therefore indicates that identity integrity should be treated as a foundational risk variable rather than simply another STRIDE category.

Session analysis produces a different risk pattern. The initial authentication event may be secure while the subsequent session becomes vulnerable. This creates a temporal separation between authentication assurance and continuing access assurance. A secure federation must therefore evaluate the complete session lifecycle rather than focusing exclusively on login security. Session-related information disclosure and unauthorized reuse can effectively transform a previously legitimate authentication event into an unauthorized access mechanism.

The access-control layer produces the most direct connection between identity and privilege. Elevation-of-privilege threats become particularly important where federated identity attributes are mapped to application-specific roles. An identity can therefore remain authentic while the resulting authorization decision is incorrect. This finding emphasizes that identity assurance and authorization correctness should be assessed as separate but dependent security properties.

The risk-oriented analysis also indicates that STRIDE categories should not be treated independently. For example, spoofing can lead to elevation of privilege, while information disclosure can facilitate session compromise. Similarly, tampering with authentication attributes can produce unauthorized access without requiring direct compromise of the identity provider.

The literature on machine learning and statistical modeling provides a useful theoretical explanation for this interconnectedness. Security risk may involve nonlinear relationships among multiple observable factors, making simple one-variable prioritization insufficient (Alzubaidi et al., 2021; Farrell et al., 2021). Statistical learning approaches likewise demonstrate the value of considering multiple explanatory variables when modeling complex outcomes (Friedman et al., 2001).

A further finding is that risk prioritization should consider propagation potential. A vulnerability affecting one isolated application may have a lower systemic risk than a weakness affecting the central identity provider or federation trust relationship. Consequently, component criticality must be incorporated into risk scoring.

Overall, the assessment indicates that the most consequential threats are those capable of crossing security boundaries: identity spoofing that reaches relying applications, authentication-data tampering that alters authorization decisions, session compromise that bypasses renewed authentication, and privilege escalation resulting from incorrect identity-to-role mapping. This reinforces the STRIDE-based federated-security perspective developed by Ganapathy (2024), while extending it through explicit risk prioritization.

DISCUSSION

The findings demonstrate that the security of federated authentication cannot be evaluated effectively through isolated examination of credentials, authentication protocols, or individual applications. Federation introduces dependencies between components, and these dependencies create systemic risks. The principal theoretical implication is that authentication security should be modeled as a chain of trust extending from identity establishment through session maintenance to authorization enforcement.

The identity layer has the greatest potential for systemic propagation. If identity trust is compromised at its source, downstream relying services may inherit the compromise. This makes spoofing and identity-data tampering especially important. The implication for security architecture is that identity assurance must remain verifiable across trust

boundaries rather than being assumed merely because the originating identity provider is considered trusted.

Session security represents a second theoretical challenge. Strong authentication does not eliminate the need to protect the resulting session state. Once authentication has succeeded, security becomes dependent on how continuing access is represented, maintained, modified, and terminated. Consequently, risk assessment should distinguish between authentication assurance and session assurance.

The third major implication concerns authorization. Federated systems can correctly authenticate a user while still granting inappropriate permissions. This occurs when identity attributes, organizational roles, or application permissions are incorrectly interpreted. Elevation of privilege therefore cannot be reduced to a conventional local-account problem; it can originate in federation-wide attribute and trust relationships.

The use of risk prioritization improves upon a purely categorical STRIDE assessment. STRIDE identifies classes of threats, but organizations must still determine which threats deserve immediate mitigation. A risk model provides this prioritization by considering likelihood, impact, component criticality, and propagation potential. Optimization research provides a conceptual basis for making decisions under constraints, while statistical modeling provides a foundation for evaluating relationships among multiple variables (Coleman and Li, 1994; Fan, 2018).

There is also a practical opportunity to integrate machine-learning techniques into future versions of the framework. Authentication systems can generate large volumes of security events, making manual interpretation increasingly difficult. Deep-learning research demonstrates the potential for representing and learning complex patterns from large datasets (Goodfellow et al., 2016; Alzubaidi et al., 2021). However, automated classification should complement rather than replace structured threat modeling. A machine-learning model may identify an anomalous event without explaining the underlying trust-boundary failure.

The principal limitation of the present study is that the supplied literature is not predominantly empirical cybersecurity literature. Several references provide mathematical, statistical, or machine-learning foundations rather than direct experimental evidence concerning federated authentication. The conclusions should therefore be interpreted as a conceptual security framework rather than statistically validated measurements of real-world attack frequency.

A second limitation is the absence of deployment-specific data. Actual risk levels vary according to architecture, authentication mechanisms, session policies, organizational trust relationships, and application-specific authorization models. The proposed framework provides a systematic assessment procedure, but quantitative risk values require empirical security-event data.

Nevertheless, the methodology has practical value because it provides a bridge between categorical threat modeling and quantitative risk prioritization. In particular, the integration of identity, session, and access analysis prevents security teams from treating authentication as a single event. Ganapathy (2024) establishes the relevance of STRIDE to federated SSO and MFA threat analysis; the present framework extends that approach toward a more explicitly risk-oriented interpretation.

CONCLUSION

Federated authentication systems provide substantial architectural benefits, but their distributed trust relationships create a security environment in which identity, session, and authorization weaknesses can propagate across multiple services. This paper presented a risk-based security assessment framework based on STRIDE, emphasizing the interconnected nature of spoofing, tampering, information disclosure, denial of service, and elevation-of-privilege threats.

The analysis identified identity establishment, session continuity, and access authorization as the three principal security domains requiring systematic assessment. Identity spoofing and assertion manipulation can undermine the foundation of federated trust; session compromise can extend unauthorized access beyond the initial authentication event; and incorrect privilege mapping can result in authorization failure even when the underlying identity is legitimate.

The major contribution of the framework is the integration of STRIDE categorization with risk prioritization. Rather than treating every threat category as equally significant, the proposed approach considers likelihood, impact, component criticality, and propagation potential. This creates a more operationally meaningful basis for deciding which weaknesses require priority mitigation.

The supplied statistical and machine-learning literature further suggests opportunities for future development.

Complex security-event datasets could be analyzed using statistical learning or optimized predictive models to supplement expert-driven STRIDE assessment. Such integration could support continuous risk scoring, anomaly detection, and adaptive prioritization while retaining an interpretable threat-model structure.

Future research should validate the proposed framework against real federated authentication deployments and develop empirical datasets covering authentication failures, session anomalies, authorization changes, and privilege transitions. Comparative evaluation of different federation architectures could further determine how trust-boundary design influences risk propagation. Ultimately, effective federated-security assessment should combine structured threat modeling, quantitative risk analysis, and evidence-based security monitoring rather than relying on any single analytical technique.

REFERENCES

1. Alzubaidi L, Zhang J, Humaidi AJ, et al. (2021). Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data*, 8(1): 53.
2. Coleman TF, Li Y (1994). On the convergence of reflective Newton methods for large-scale nonlinear minimization subject to bounds. *Mathematical Programming*, 67(2): 189–224.
3. Fan J (2018). *Local Polynomial Modelling and Its Applications*, Monographs on Statistics and Applied Probability 66. Routledge.
4. Farrell MH, Liang T, Misra S (2021). Deep neural networks for estimation and inference. *Econometrica*, 89(1): 181–213.
5. Friedman JH (1991). Multivariate adaptive regression splines. *The Annals of Statistics*, 19(1): 1–67.
6. Friedman JH, Tibshirani R, Hastie T (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 1st ed. Springer Series in Statistics. Springer, New York.
7. 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-ssso-mfa>
8. Goodfellow I, Bengio Y, Courville A (2016). *Deep Learning*. MIT Press.