

Attack Surface Analysis of Federated SSO and Multi-Factor Authentication Systems Using the STRIDE Framework

Hiroshi Nakamura

Department of Data Science and Intelligent Systems,
Advanced Computing Research Institute, Tokyo, Japan

ARTICLE INFO

Article history:

Submission: July 08, 2026

Accepted: August 18, 2026

Published: September 01, 2026

VOLUME: Vol.11 Issue 09 2026

Keywords:

Federated SSO, Multi-Factor Authentication, STRIDE, Threat Modeling, Attack Surface, Identity Federation, Authentication Security, Access Control, Cybersecurity

ABSTRACT

Federated Single Sign-On (SSO) and Multi-Factor Authentication (MFA) systems have become fundamental components of modern identity and access management architectures because they reduce credential duplication while enabling centralized authentication across multiple applications and organizational domains. However, federation introduces complex trust relationships among identity providers, service providers, authentication protocols, tokens, sessions, and user devices, creating an attack surface that cannot be adequately assessed by examining individual authentication components in isolation. This research examines the attack surface of federated SSO and MFA systems through a STRIDE-oriented threat modeling perspective. The methodology decomposes the authentication ecosystem into identity, protocol, token, session, endpoint, and trust-boundary components and evaluates potential threats involving spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege. The analytical approach is informed by the supplied literature on feature selection, sparse modeling, classification, statistical discrimination, and high-dimensional data analysis, which provides theoretical foundations for reducing complex threat spaces into interpretable risk dimensions. The analysis demonstrates that federation can simultaneously improve centralized security governance and increase systemic dependency on identity infrastructure. MFA reduces several credential-related risks but does not eliminate attacks against authentication workflows, tokens, sessions, recovery mechanisms, or federation trust relationships. The study proposes a structured attack-surface model in which threat prioritization is based on component criticality, trust-boundary exposure, attack-path concentration, and potential privilege impact. The findings emphasize that effective federated authentication security requires security assessment beyond login mechanisms and must incorporate the complete identity transaction lifecycle.

INTRODUCTION

Modern enterprise environments increasingly depend on distributed applications, cloud services, partner platforms, and heterogeneous computing environments. In such environments, users frequently require access to numerous applications while organizations seek to maintain centralized control over authentication and authorization. Federated Single Sign-On (SSO) addresses this requirement by allowing an identity provider to authenticate a user and communicate authentication assertions or tokens to relying service providers. Multi-Factor Authentication (MFA) strengthens this process by requiring multiple categories of evidence rather than relying exclusively on passwords.

Despite these advantages, federated authentication changes the security architecture from an isolated login model into a network of interdependent trust relationships. A compromise of an identity provider, federation configuration, authentication token, or privileged administrative interface may affect multiple downstream services simultaneously. Consequently, the security of a federated ecosystem depends not only on the strength of individual authentication factors but also on the integrity of the interactions connecting identity providers, service providers, users, tokens, sessions, and recovery mechanisms.

Threat modeling provides a systematic approach for identifying and evaluating such risks. In particular, the STRIDE framework provides six threat categories: spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege. Applying these categories to federated authentication enables security analysts to examine both direct vulnerabilities and attack chains that cross multiple trust boundaries. Ganapathy (2024) specifically positions STRIDE-based analysis as a useful approach for identifying attack vectors within federated SSO and MFA systems, providing the central conceptual basis for this study.

The methodological challenge is that a federated identity architecture contains numerous interacting variables. Conventional qualitative assessment can therefore become difficult when many services, authentication factors, protocols, identities, and privilege levels are involved. The supplied literature on feature selection and dimensionality reduction is relevant to this challenge because it demonstrates how complex, high-dimensional data can be reduced to meaningful variables without losing critical analytical structure (Boulesteix and Strimmer, 2006; Chun and Keleş, 2010). Similarly, classification research provides useful foundations for distinguishing different categories of observations and identifying discriminative characteristics (Chung and Keles, 2010; Freeman et al., 2013).

The objective of this research is therefore to construct a systematic attack-surface analysis model for federated SSO and MFA systems using STRIDE while drawing methodological support from the supplied literature on classification, variable selection, dimensionality reduction, and statistical modeling. The research focuses on identifying security-relevant components, mapping trust boundaries, classifying threats, prioritizing attack surfaces, and interpreting the resulting security implications.

LITERATURE REVIEW

The supplied literature does not directly form a homogeneous cybersecurity literature base; instead, much of it addresses statistical learning, feature selection, classification, dimensionality reduction, and pattern recognition. These methodological contributions are nevertheless relevant to threat modeling because an enterprise authentication ecosystem can be represented as a complex set of components and observable security characteristics.

Boulesteix and Strimmer (2006) describe partial least squares as a flexible technique for analyzing high-dimensional genomic data. The underlying methodological insight is important for threat modeling: when a system contains numerous correlated variables, meaningful latent structures can be extracted to support analysis. Federated authentication similarly contains correlated security variables, including identity provider exposure, authentication factor strength, token characteristics, session duration, application privilege, and trust relationships.

Chun and Keleş (2010) develop sparse partial least squares regression for simultaneous dimension reduction and variable selection. This is particularly relevant to attack-surface analysis because not every observable system characteristic contributes equally to security risk. A useful threat model must distinguish high-impact attack-surface variables from relatively insignificant attributes. Sparse modeling provides a conceptual basis for retaining informative variables while reducing analytical complexity.

Chung and Keles (2010) extend sparse partial least squares to classification. Their work is useful for considering the categorization of security observations into meaningful threat classes. Within STRIDE, individual attack conditions can be categorized according to threat type, affected component, trust boundary, and privilege consequence.

De Jong (1993) presents SIMPLS as an alternative approach to partial least squares regression. The broader contribution is methodological: structured transformation of complex correlated variables can make analytical relationships more manageable. In federated authentication, multiple security indicators may represent overlapping aspects of the same underlying weakness. For example, insecure token handling can simultaneously contribute to spoofing, tampering, information disclosure, and elevation of privilege.

The variable-selection literature further supports risk prioritization. Fan and Li (2001) demonstrate the importance of variable selection under nonconcave penalization, while Fan et al. (2009) examine feature selection in ultrahigh-dimensional environments. These perspectives suggest that threat modeling should not treat every possible attack vector as equally significant. Instead, analysts should identify the variables that most strongly differentiate high-risk attack paths from lower-impact conditions.

Freeman et al. (2013) examine feature-selected tree-based classification, offering another conceptual foundation for hierarchical threat classification. A federated authentication system can similarly be represented as a hierarchical decision structure in which an attack is first associated with an architectural layer, then a component, then a STRIDE category, and finally an impact class.

Friedman et al. (2010) discuss regularization paths for generalized linear models, reinforcing the broader principle that model complexity should be controlled when multiple predictive variables exist. In threat modeling, uncontrolled enumeration can generate excessively large threat catalogs that reduce practical usefulness. Structured prioritization is therefore essential.

The statistical classification foundations of Fisher (1936) and McLachlan (2005) further demonstrate the importance of discriminating observations according to relevant characteristics. This principle translates naturally to cybersecurity risk assessment, where threats must be distinguished according to attack mechanism, affected asset, likelihood, and consequence.

Johnstone and Silverman (2004) address sparse sequences and empirical Bayes estimation, while Loh (2011) discusses classification and regression trees. Together, these studies reinforce the value of identifying informative variables and hierarchical decision boundaries.

The supplied biomedical and genomic references, including Hutter and Zenklusen (2018), Nguyen and Rocke (2002a, 2002b), and Tam et al. (2019), demonstrate how large, heterogeneous datasets can contain valuable relationships requiring structured analysis. Although their application domains differ from identity security, their methodological relevance lies in handling complex feature spaces.

The central literature gap is therefore not the absence of statistical methodologies but the limited direct integration of such analytical principles with federated authentication threat modeling. Ganapathy (2024) provides a direct foundation for STRIDE-based analysis of federated SSO and MFA, while the other supplied works contribute methodological perspectives for selecting, reducing, and classifying complex security variables. This study positions these perspectives within an integrated attack-surface framework.

METHODOLOGY

3.1 Research Design

This research adopts a conceptual and analytical threat-modeling methodology. The approach does not attempt to measure a specific organization's infrastructure; instead, it develops a generalized analytical model applicable to federated SSO and MFA architectures.

The methodology consists of five principal stages: system decomposition, trust-boundary identification, attack-surface extraction, STRIDE threat classification, and risk prioritization. Each stage transforms architectural complexity into a structured analytical representation.

The conceptual foundation follows the principle that a complex system should first be represented through meaningful variables before classification. This principle is consistent with research on dimensionality reduction and feature selection (Boulesteix and Strimmer, 2006; Chun and Keleş, 2010).

3.2 Federated Authentication System Decomposition

A federated SSO and MFA environment can be divided into six analytical domains: the user endpoint, identity provider, authentication-factor infrastructure, federation protocol layer, service provider, and administrative/control plane.

The user endpoint includes browsers, mobile devices, workstations, and associated authentication applications. Its attack surface includes credential theft, session compromise, malicious extensions, device manipulation, and interception of authentication workflows.

The identity provider (IdP) is a high-value security component because it performs or coordinates user authentication and issues identity assertions. A successful compromise of the IdP may have a disproportionate effect because the attacker can potentially gain access to multiple relying applications.

The MFA layer includes passwords, one-time codes, authenticator applications, hardware security mechanisms, push-based authentication, and recovery processes. MFA reduces the risk associated with stolen passwords, but its security depends on the resistance of the entire authentication process rather than merely the existence of a second factor.

The federation protocol layer transports identity information between trusted entities. Tokens or assertions become security-critical objects because their integrity and validity determine whether service providers accept an authentication event.

The service-provider layer contains applications that consume federated identities. Weak authorization controls at this level can transform valid authentication into unauthorized access.

Finally, the administrative plane contains federation configuration, identity lifecycle management, role assignments, application registration, trust relationships, and security policies. Administrative compromise can create systemic consequences even when user authentication itself remains technically strong.

3.3 Trust-Boundary Analysis

Trust boundaries identify transitions where information, authentication assertions, or privileges move between security domains. In federated SSO, important boundaries exist between the user and IdP, IdP and MFA services, IdP and service provider, browser and application, and administrators and identity infrastructure.

A hypothetical enterprise may operate one central IdP supporting twenty applications. If one application is compromised, the potential impact may remain localized. If the IdP or federation trust configuration is compromised, however, the attacker may influence authentication decisions for all twenty applications. Therefore, attack-surface analysis must consider not only individual vulnerability severity but also the concentration of trust.

This concept is consistent with feature-selection reasoning: variables associated with broader system behavior should receive greater analytical significance than variables with narrow effects (Fan et al., 2009).

3.4 STRIDE Threat Mapping

The STRIDE framework categorizes threats into six classes.

Spoofing concerns unauthorized impersonation. In federated systems, spoofing can involve compromised credentials, stolen authentication tokens, fraudulent identity assertions, or compromised sessions. MFA can reduce password-based spoofing but cannot automatically prevent every form of identity impersonation.

Tampering concerns unauthorized modification of data or security-relevant transactions. In a federated architecture, token manipulation, altered federation configuration, or modification of authorization attributes can potentially produce unauthorized outcomes.

Repudiation concerns the inability to reliably establish responsibility for an action. Authentication logs, token issuance records, administrative changes, and session events are therefore important evidence sources.

Information disclosure concerns unauthorized exposure of credentials, authentication tokens, identity attributes, session information, or sensitive configuration data.

Denial of service addresses attacks that prevent legitimate authentication or application access. Because SSO centralizes authentication, disruption of an IdP can potentially affect multiple applications simultaneously.

Elevation of privilege occurs when an attacker moves from a lower privilege state to a higher one. This can arise through authorization mapping errors, manipulated identity attributes, compromised administrative roles, or weaknesses in application-to-IdP trust.

The STRIDE mapping follows the central premise that threat categories should be systematically associated with system components and attack paths rather than considered independently (Ganapathy, 2024).

3.5 Attack-Surface Feature Model

To prioritize threats, this research defines a conceptual feature vector for each attack surface:

$$A = \{E, T, P, I, C, R\}$$

where E represents exposure, T represents trust-boundary sensitivity, P represents privilege impact, I represents information sensitivity, C represents centrality, and R represents recoverability.

Exposure measures how accessible a component is to potential attackers. Trust-boundary sensitivity represents the number and importance of security-domain transitions associated with the component. Privilege impact measures the potential increase in attacker authority. Information sensitivity reflects the value of accessible information. Centrality represents the number of dependent applications or authentication flows. Recoverability considers how quickly security operations can restore trustworthy service.

The model is deliberately conceptual rather than numerical. Its purpose is to demonstrate how high-dimensional architectural observations can be reduced into security-relevant variables. This approach is consistent with the methodological rationale of sparse variable selection and dimensionality reduction (Chun and Keleş, 2010; De Jong, 1993).

3.6 Threat Prioritization

Threat prioritization should consider both local vulnerability characteristics and systemic consequences. For example, a low-complexity weakness in an individual application may be less consequential than a moderate weakness in the central IdP because the latter can influence numerous dependent services.

A classification-oriented approach can therefore assign threats to priority classes such as critical, high, moderate, and low according to the combination of exposure, privilege impact, trust centrality, and affected information. Classification and regression methods provide conceptual support for this type of structured differentiation (Loh, 2011; McLachlan, 2005).

A hypothetical example illustrates the approach. Suppose an enterprise has an IdP connected to twenty-five applications, each with different authorization roles. A stolen password protected by MFA may initially represent a limited threat. However, if the recovery mechanism allows an attacker to bypass strong authentication, the risk becomes substantially greater because the compromised identity can potentially reach multiple applications. The attack surface is therefore the entire recovery workflow, not simply the password.

3.7 Validation Logic

Because this study is conceptual, validation focuses on analytical consistency rather than empirical penetration testing. Each identified threat should satisfy four conditions: it must affect an identifiable component, cross or depend upon a meaningful trust boundary, map to at least one STRIDE category, and produce a plausible security consequence.

The feature-selection literature supports this principle by emphasizing the value of retaining informative variables while reducing unnecessary complexity (Fan and Li, 2001; Fan et al., 2009). A threat model that generates hundreds of theoretically possible threats without prioritization may become operationally ineffective.

RESULTS

The analysis indicates that the federated SSO and MFA attack surface is structurally concentrated around a relatively small number of high-value components. The identity provider, federation trust configuration, authentication-token lifecycle, MFA recovery process, privileged administrative interfaces, and session-management mechanisms represent the most consequential areas because compromise of these components can influence multiple stages of the identity transaction.

The first major finding is that centralization creates both security efficiency and systemic dependency. SSO reduces the number of credentials users must manage and enables centralized authentication policy. However, the same centralization increases the potential blast radius of an IdP compromise. The identity provider should consequently be modeled as a high-centrality component rather than simply another authentication endpoint.

The second finding is that MFA changes rather than eliminates the attack surface. Strong authentication factors can substantially improve resistance against password-only compromise, but attackers may instead target session tokens, recovery processes, device trust, authorization mappings, or federation configuration. Consequently, security analysis limited to MFA-factor strength produces an incomplete assessment.

The third finding is that trust boundaries are more informative than component inventories alone. Two components with similar technical functionality can have very different security importance depending on the number of trust relationships they maintain. A service provider with limited federation dependency has a narrower systemic impact than a central identity component supporting numerous applications.

The fourth finding concerns attack-path convergence. Multiple STRIDE categories can emerge from a single underlying weakness. For example, compromised federation configuration may facilitate tampering, information disclosure, spoofing, and elevation of privilege simultaneously. This means threat modeling should analyze attack chains rather than assign exactly one threat category to every vulnerability.

The fifth finding is that feature selection is useful for reducing threat-model complexity. The statistical literature demonstrates that high-dimensional problems can contain a smaller set of informative variables (Boulesteix and Strimmer, 2006; Chun and Keleş, 2010). Applied conceptually to federated authentication, exposure, privilege impact, centrality, trust sensitivity, and information sensitivity can provide a compact representation of attack-surface importance.

The sixth finding is that authorization must be evaluated together with authentication. A federated system may correctly establish identity while still assigning excessive application privileges. Therefore, a secure authentication architecture cannot be evaluated exclusively through successful login protection.

Finally, the analysis confirms the importance of systematic STRIDE mapping for federated authentication. Ganapathy (2024) emphasizes the usefulness of STRIDE-based attack-vector analysis for federated SSO and MFA systems. The present model extends that perspective by emphasizing feature reduction, attack-surface centrality, and structured prioritization.

DISCUSSION

The findings demonstrate that the security of federated SSO and MFA should be understood as an ecosystem property rather than a property of individual authentication mechanisms. The principal theoretical implication is that authentication strength and architectural resilience represent different dimensions of security. MFA may increase authentication assurance, while federation architecture determines how far the consequences of a successful compromise can propagate.

The distinction is important because organizations may incorrectly interpret the deployment of MFA as evidence that their identity infrastructure has become comprehensively secure. MFA addresses particular authentication weaknesses, especially those associated with password compromise, but an attacker may target another component in the authentication lifecycle. Token theft, session hijacking, recovery abuse, configuration compromise, and excessive authorization can remain relevant attack surfaces.

The STRIDE model provides a useful conceptual structure because it prevents threat analysis from becoming limited to credential theft. Spoofing identifies identity impersonation; tampering captures manipulation of security information; repudiation emphasizes evidence and accountability; information disclosure addresses confidentiality; denial of service addresses availability; and elevation of privilege addresses unauthorized authority. Ganapathy (2024) demonstrates the relevance of this approach specifically to federated SSO and MFA environments.

The literature on statistical classification and variable selection provides a complementary perspective. Threat models can suffer from excessive dimensionality because modern identity environments contain large numbers of applications, identities, policies, devices, tokens, roles, and configurations. Feature-selection approaches suggest that analysts should identify the variables that provide the greatest discriminatory value when distinguishing high-risk from lower-risk conditions (Fan and Li, 2001; Fan et al., 2009). Similarly, classification approaches provide a theoretical basis for grouping threats into interpretable categories (Chung and Keles, 2010; McLachlan, 2005).

A significant trade-off concerns centralization. Centralized authentication improves policy consistency and reduces duplicated credential-management processes. At the same time, it produces high-value infrastructure whose disruption or compromise may have widespread consequences. Thus, security controls for central identity services should be proportionate to their systemic centrality.

Another trade-off concerns usability. Additional authentication factors and security checks can reduce attack probability but may introduce user friction and increase support requirements. Recovery mechanisms illustrate this tension particularly clearly: overly permissive recovery can weaken authentication, while excessively restrictive recovery can create operational problems. Therefore, recovery must be included within the threat model as an authentication-equivalent pathway.

The analysis also reveals a limitation in applying statistical feature-selection concepts directly to cybersecurity. Security threats are adversarial and adaptive, whereas many statistical modeling environments assume relatively stable relationships among variables. Attackers can deliberately change behavior to evade classification. Consequently, statistical techniques should inform analytical organization rather than be interpreted as definitive predictors of future attacks.

Another limitation is the conceptual nature of this study. No empirical organizational dataset, penetration-testing results, incident records, or quantitative attack probabilities were used. The findings therefore establish a structured analytical framework rather than statistically validated risk probabilities. Future empirical studies could evaluate the framework against real federated infrastructures and compare threat-priority rankings across different architectures.

Nevertheless, the combined framework has practical significance. Security teams can use architectural decomposition to identify high-centrality components, STRIDE to classify threats, and feature-selection principles to reduce analytical complexity. Such integration can make threat modeling more actionable by focusing assessment resources on attack surfaces that have both high exposure and high systemic impact.

CONCLUSION

Federated SSO and MFA systems provide substantial architectural benefits by centralizing authentication, reducing credential duplication, and supporting secure access across distributed applications. However, federation also creates complex trust relationships and concentrates security dependencies within identity infrastructure.

This research developed a STRIDE-oriented attack-surface analysis model that decomposes federated authentication into user endpoints, identity providers, MFA mechanisms, federation protocols, service providers, and administrative infrastructure. The analysis demonstrates that identity providers, federation trust configurations, authentication-token lifecycles, recovery mechanisms, sessions, and authorization mappings represent particularly significant attack surfaces.

The research also integrates methodological principles from the supplied literature on dimensionality reduction, sparse variable selection, classification, and statistical discrimination. These principles support a structured approach in which large numbers of architectural variables are reduced to meaningful security characteristics such as exposure, trust sensitivity, privilege impact, information sensitivity, centrality, and recoverability.

The principal contribution is the positioning of federated authentication security as a combined problem of authentication assurance, trust management, attack-surface centrality, and authorization integrity. MFA should therefore not be treated as the final security boundary. Instead, security assessment must encompass the complete identity transaction lifecycle.

Future research should empirically validate the proposed feature model using real-world federated architectures, security-event datasets, penetration-testing observations, and longitudinal incident data. Quantitative scoring models could subsequently be developed to estimate the relative contribution of different attack-surface features. Such research could further connect threat modeling with automated security analytics while retaining the interpretability required for enterprise security governance.

REFERENCES

1. Boulesteix AL, Strimmer K (2006). Partial least squares: A versatile tool for the analysis of high-dimensional genomic data. *Briefings in Bioinformatics*, 8: 32–44.
2. Chun H, Keleş S (2010). Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 72(1): 3–25.
3. Chung D, Keles S (2010). Sparse partial least squares classification for high dimensional data. *Statistical Applications in Genetics and Molecular Biology*, 9. Article 17.
4. De Jong S (1993). Simpls: An alternative approach to partial least squares regression. *Chemometrics and Intelligent Laboratory Systems*, 18: 251–263.
5. Fan J, Li R (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96: 1348–1360.
6. Fan J, Samworth R, Wu Y (2009). Ultrahigh dimensional feature selection: Beyond the linear model. *Journal of Machine Learning Research*, 10: 2013–2038.
7. Fanty M, Cole R (1990). Spoken letter recognition. *Proceedings of the International Conference on Neural Information Processing Systems*, 4: 220–226.
8. Fisher RA (1936). The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2): 179–188.

9. Freeman C, Kulić D, Basir O (2013). Feature-selected tree-based classification. *IEEE Transactions on Cybernetics*, 43(6): 1990–2004.
10. Friedman J, Hastie T, Tibshirani R (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1): 1.
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.
12. Hoskuldsson A (1988). PLS regression methods. *Journal of Chemometrics*, 2: 211–228.
13. Hoskuldsson A (1992). The h-principle in modelling with applications to chemometrics. *Chemometrics and Intelligent Laboratory Systems*, 14: 139–153.
14. Hutter C, Zenklusen JC (2018). The Cancer Genome Atlas: Creating lasting value beyond its data. *Cell*, 173(2): 283–285.
15. Johnstone IM, Silverman BW (2004). Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences. *The Annals of Statistics*, 32(4): 1594–1649.
16. Lê Cao KA, Rossouw D, Robert-Granié C, Besse P (2008). A sparse PLS for variable selection when integrating omics data. *Statistical Applications in Genetics and Molecular Biology*. 7(1): Article 35.
17. Lin Y, Zhang M, Zhang D (2015). Generalized orthogonal components regression for high dimensional generalized linear models. *Computational Statistics & Data Analysis*, 88: 119–127.
18. Loh WY (2011). Classification and regression trees. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1): 14–23.
19. Massy WF (1965). Principal components regression in exploratory statistical research. *Journal of the American Statistical Association*, 60(309): 234–256.
20. McLachlan GJ (2005). *Discriminant Analysis and Statistical Pattern Recognition*. John Wiley & Sons.
21. Nguyen DV, Rocke DM (2002a). Classification of Acute Leukemia Based on DNA Microarray Gene Expressions Using Partial Least Squares. Springer.
22. Nguyen DV, Rocke DM (2002b). Tumor classification by partial least squares using microarray gene expression data. *Bioinformatics*, 18: 39–50.
23. Shen J, Gao S (2008). A solution to separation and multicollinearity in multiple logistic regression. *Journal of Data Science*, 6(4): 515.
24. Tam V, Patel N, Turcotte M, Bossé Y, Paré G, Meyre D (2019). Benefits and limitations of genome-wide association studies. *Nature Reviews. Genetics*, 20(8): 467–484.
25. 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-ssomfa>