ADAPTIVE MULTI-FACTOR AUTHENTICATION IN A DOCUMENT-CENTRIC SYSTEM

UDC: 004.056.523
DOI: https://doi.org/10.53486/tids2025.09

GHENADIE BELINSCHI
Information Security Laboratory of the Academy of Economics Studies,
ghenadie.belinschi@ase.md
ORCID ID: 0009-0009-3361-9890

Abstract: Static authentication methods are increasingly vulnerable to social engineering and phishing, especially in cloud-based document workflows. This paper presents an adaptive multi-factor authentication model that escalates security measures with higher document criticality or anomalous user behavior. By merging biometric data with cryptographic safeguards under Zero-Trust principles, continuous verification is achieved while balancing usability. Preliminary findings suggest a notable increase in protection for sensitive documents and underscore the need to standardize event-driven controls.

Key words: Multi-factor authentication, document-centric systems, data protection.

JEL Classification: C 88, O 33, M 15

INTRODUCTION

In the era of digital transformation, the need for reliable user authentication when handling sensitive data is steadily increasing. Traditional authentication methods—such as relying solely on passwords—are no longer sufficient to ensure security. Passwords can be easily guessed or stolen through phishing and data breaches. According to Verizon, approximately one-quarter of all data breaches involve the use of stolen credentials (Verizon, 2022). Even the implementation of static two-factor authentication (2FA) does not fully resolve the issue: social engineering and phishing attacks can still bypass one-time codes, while the constant demand for additional authentication steps significantly reduces user convenience.

The Zero Trust concept calls for a fundamental shift in how access security is approached. At its core is the principle of "never trust, always verify," meaning that every attempt to access a resource is treated as potentially unsafe until proven otherwise (NIST, 2020). Traditional perimeter-based security models assume that internal users and devices can be trusted by default. In contrast, Zero Trust is built on a presumption of distrust: authentication and authorization must be performed for every access request, regardless of whether it originates inside the corporate network or from outside (NIST, 2020). In practice, Zero Trust architecture requires continuous verification of both user and device during every session or transaction, along with enforcing the minimum necessary access level for each request. This approach significantly reduces risks associated with both external attacks and internal threats, including compromised employee credentials or devices.

Nonetheless, applying Zero Trust principles introduces a challenge to user convenience: if all possible authentication factors are requested for every action, the user experience becomes excessively burdensome. Therefore, a key objective is to strike a balance between stringent access control and usability. One effective approach is adaptive multi-factor authentication—a dynamic method where the system tailors verification procedures according to the current risk level. In an adaptive model, low-risk actions are executed with little to no disruption for the user, whereas higher-

risk situations trigger requests for additional identity verification. This risk-based strategy has gained significant attention from experts, as it addresses the long-standing trade-off between security and convenience (Brodsky, A., 2018). Specifically, risk-based authentication (RBA) determines not only whether authentication is required, but also which and how many factors should be used, based on a real-time assessment of the current context (Brodsky, A., 2018).

Document-centric systems have specific authentication requirements. Access to protected documents must be granted only to verified users with the appropriate permissions, and each operation—such as opening, editing, or sending a document—should be explicitly confirmed. At the same time, employees require convenient and seamless access to documents in order to work efficiently. This creates a demand for adaptive verification mechanisms that can strengthen security in response to abnormal or high-risk activities, while remaining unobtrusive during routine, low-risk actions.

The aim of this study is to design and analyze an architecture for adaptive multi-factor authentication in a document-centric system, aligned with the principles of Zero Trust.

1. METHODOLOGY

This study is applied in nature and follows the design science methodology commonly used in information security research. It begins with an analysis of the shortcomings of traditional authentication mechanisms and the requirements imposed by the Zero Trust concept. Based on this analysis, the architecture of an adaptive multi-factor authentication (MFA) system tailored for a document-centric environment was developed. The system's core components were implemented in a prototype and tested using model scenarios involving access to protected documents. To evaluate the outcomes, a combination of comparative analysis and experimental simulation of access attempts—both legitimate and malicious—was employed.

The authentication system comprises several modules that interact through an event-driven model. The key components of the architecture and their roles are as follows:

The authentication module serves as the central component responsible for processing access requests to documents. Each time a user attempts to access a protected document, this module initiates a verification process—prompting the user to provide the necessary authentication factors based on the assessed risk level. It supports a range of factors, including passwords, one-time passcodes (OTP), biometric data (such as fingerprints or facial recognition), and hardware tokens. Under normal, low-risk conditions, a single-factor check (e.g., password) may be sufficient. However, when the risk level increases, additional factors are automatically enforced. In this way, the module implements the Step-Up authentication approach, escalating the verification requirements as needed.

The Risk Engine is a dedicated module for risk assessment that analyzes the context of each session and assigns it a dynamic risk level. It collects input data from multiple sources—such as user identity, device characteristics, network details, time of access, and the type of operation being requested. Using this data, it calculates a real-time risk score based on predefined rules and models. Relevant risk factors include the user's geographic location and deviations from their typical patterns, whether the device and browser are recognized or new, the time of access (e.g., during or outside working hours), and whether the user's current behavior aligns with their historical activity patterns (Verizon, 2022), (Brodsky, A., 2018).

The Risk Engine can also leverage behavioral biometrics, such as typing speed and rhythm or mouse movement patterns, to help identify users by their unique behavior. Based on the combination

of these inputs, the system classifies each session as low, medium, or high risk. The authentication module then adjusts its response accordingly: low-risk sessions proceed transparently, medium-risk sessions require an additional verification factor, and high-risk sessions are blocked—prompting extended checks or outright denial. This adaptive strategy enables the system to enhance security without disrupting the experience for legitimate users.

The event model is designed to be event-driven. This means that all significant events and context changes are captured and formalized as structured records, which are then transmitted to the Risk Engine and other system modules. The model encompasses various types of events, including authentication events (e.g., successful login, failed login attempt, additional factor request), application events (e.g., document opening, file upload or transmission), and system events (e.g., session IP address change, new device connection, privilege escalation).

Each event is described using a defined set of attributes—such as event type, timestamp, user ID, and contextual data like geolocation or device status. This standardized event format ensures consistent and reliable information flow into the Risk Engine. The model builds on established approaches in security event management and user and entity behavior analytics (UEBA). However, the Zero Trust paradigm currently lacks a unified standard for event modeling, making the integration of disparate systems more complex (Morrow *et al.*, 2022). To address this, our solution uses a custom event format and vocabulary, which may serve as a foundation for future standardization efforts.

Cryptographic Protection. Confidential user data—such as password hashes and biometric templates—is stored in the authentication database in encrypted form. Robust, standardized encryption algorithms are used to ensure data security (e.g., AES-256 for data storage and TLS 1.3 for data transmission), meeting the requirements for confidentiality and integrity.

Additional safeguards are implemented to prevent the compromise of authenticators themselves. For example, biometric data is transmitted with a cryptographic signature that verifies both its authenticity and its origin from a trusted sensor. When hardware tokens are used (e.g., smart cards or USB keys), their private keys never leave the device; instead, the token signs authentication challenges internally, eliminating the risk of secret interception. In this way, the architecture adheres to Zero Trust principles by securing all communications—regardless of network location—and validating the authenticity of data sources (NIST, 2020).

Sequence diagrams and use case models were developed to illustrate the system architecture and the interaction between its components. The paper also includes pseudocode for the Risk Engine rules and examples of event formats to demonstrate the system's logic. Following the architectural design, a prototype was implemented: a web-based document management service integrated with the authentication module and a simplified version of the Risk Engine. This prototype enabled simulation testing, the results of which are discussed in the following section.

2. COMPARISON OF TRADITIONAL AND ADAPTIVE AUTHENTICATION SCHEMES

To assess the effectiveness of the proposed system, we conducted a comparative analysis against traditional approaches. In a document-centric environment, traditional authentication typically relies on static multi-factor authentication (MFA)—for example, a user is always required to enter a password followed by a one-time code sent to their phone upon login. While this method offers better security than password-only access, it has notable limitations.

First, it lacks contextual awareness—additional factors are always required, even in low-risk situations (such as when an employee logs in from their usual workstation in the office). Second,

static MFA does not respond to changes in real time. Once the initial check is passed, access remains open, and the system does not continue to monitor or evaluate subsequent activity.

The proposed adaptive authentication approach is characterized by the following key features.

Context-awareness – The decision to require MFA is based on real-time contextual information about the session. The system evaluates factors such as the device and location from which the request originates, the type of document being accessed, the time of day, and more. If no anomalies are detected, additional authentication steps are skipped, enhancing the overall user experience (Brodsky, A., 2018).

Dynamic and continuous – Risk assessment is performed for every new access event. Under the Zero Trust model, each action must be verified independently, meaning that even after a successful login, the system may require re-authentication when accessing highly sensitive documents or when contextual changes occur—such as a sudden change in the session's IP address (NIST, 2020). In contrast, traditional models often grant broad access after the initial login, which poses a risk if the session is compromised. The adaptive model follows a "per request" decision-making principle—each request is evaluated in real time to determine whether additional verification is necessary.

Balance of security and convenience – The system is designed to deliver strong security where it is necessary, without overburdening users in low-risk scenarios. As noted, risk-based authentication helps reconcile the tension between security and user convenience (Brodsky, A., 2018). In typical conditions, users barely notice the protection mechanisms—login is seamless, and access to standard documents requires no extra steps. In contrast, when the situation is atypical, the system escalates the level of verification. According to recent research, users find adaptive (risk-based) authentication more convenient than constant two-factor authentication, while also perceiving it as more secure than password-only access (Wiefling et al., 2021). As a result, the proposed approach improves user satisfaction without compromising security.

Resistance to credential compromise – If an attacker obtains a user's password, traditional MFA may still prevent access—provided the second factor is enabled. However, many users disable the second factor due to inconvenience, or use weaker methods such as SMS, which can be intercepted. In an adaptive scheme, the compromise of a password alone is not enough to breach the system. If a login attempt is made from an unfamiliar device or an unusual location, the Risk Engine assigns a high risk level and prompts for additional factors that the attacker likely cannot provide (such as biometrics or a hardware token). The system can also fully block suspicious attempts if the assessed likelihood of an attack is high (Brodsky, A., 2018). This significantly reduces the risk of account takeover through stolen credentials.

To illustrate how the scheme functions, consider the following example scenarios.

Scenario 1 (typical access): An employee logs into a document-centric system from their office computer. The device is registered, the geolocation matches the office location, and the system recognizes that the user has previously logged in from this device. In addition, behavioral indicators—such as typing speed—match the user's typical pattern. The Risk Engine classifies the session as low risk. As a result, only basic authentication is required—for example, entering a password—after which the user is granted immediate access to documents without additional checks. From the user's perspective, this feels like a simple, seamless login.

Scenario 2 (high risk): The same employee attempts to log in from an internet café in another country and requests access to a financial report marked "Confidential." In this case, the context is unusual: the device is unknown, the location is atypical, and the requested document is highly

sensitive. The Risk Engine assigns a high risk level to the session. Consequently, the system may require multiple authentication factors: in addition to the password, the user must confirm the login via a mobile app (e.g., push notification or OTP) and complete biometric verification. Access is granted only if all required factors are successfully provided. If the attempt is fraudulent—such as when an attacker cannot supply valid biometrics or the correct OTP—access is denied, and the incident is forwarded to the security monitoring system. In this way, the adaptive mechanism responds intelligently to context: in the first case, the user experiences no friction, while in the second, a potential attack is thwarted through multi-layered verification.

3. SYSTEM SECURITY AND THREAT MANAGEMENT

Security Analysis of Adaptive MFA – The proposed system offers a significantly higher level of protection compared to static authentication schemes. Even if an attacker manages to bypass one factor—such as by guessing or stealing a password—the likelihood of simultaneously overcoming multiple independent factors, while also passing behavioral and contextual checks, is extremely low. Particular emphasis is placed on resilience against common types of attacks.

As noted earlier, a password alone does not guarantee secure access—phishing and password guessing remain significant threats. In cases of suspicious login attempts, the system requires additional authentication factors that phishing sites cannot easily intercept, such as biometric data or a push notification to a trusted device. Moreover, the Risk Engine can detect anomalies characteristic of automated attacks—for instance, unusually high-speed credential entry or login attempts using bulk username lists—and immediately flag such sessions as high-risk, thereby blocking mass brute-force attempts.

In traditional models, if an attacker gains access to an active session cookie, they can often continue interacting with the system without further authentication. In contrast, the adaptive model enforces re-authentication for every critical action. For example, even with a stolen session cookie, an attacker attempting to open a protected document or perform actions on behalf of the user would be prompted for additional authentication (re-authentication). Furthermore, the system can detect if the session has been transferred to a different device or IP address, flag it as potentially hijacked, and require the user to re-authenticate entirely.

A dishonest employee or an external attacker who gains access to the internal network cannot move freely across resources without oversight. In a Zero Trust model, every action across different nodes requires explicit authorization (NIST, 2020). To mitigate internal threats and privilege escalation, the system logs all document access requests and can identify suspicious patterns of behavior—even from users who are legitimately logged in. For instance, opening a large number of documents in rapid succession or attempting access at unusual hours may trigger a higher risk classification. As a result, the system may prompt the user to re-authenticate or temporarily suspend access pending further investigation. In this way, the system hinders insider attacks and lateral movement within the network.

An important component of the system is the use of biometrics—such as fingerprint, facial, or voice recognition—as one of the authentication factors. Biometric authenticators offer a significant advantage: they cannot be forgotten or shared, are inherently linked to the individual user, and—unlike passwords or tokens—cannot be accidentally lost. However, biometrics also come with known vulnerabilities. One issue involves identification errors (false acceptances and false rejections), where the system may confuse one user for another. While modern algorithms minimize these errors, a more

serious concern is spoofing—i.e., the forgery of biometric data. An attacker might try to trick the system using a fake fingerprint or a photograph of a face instead of a live person. It is well-documented that simple techniques, like showing a printed photo to a camera, can deceive some facial recognition systems if no additional checks are in place (Zakuanova *et al*., 2018).

To counter such threats, the system incorporates anti-spoofing mechanisms, collectively referred to as Presentation Attack Detection (PAD). The first layer of protection involves a liveness check during biometric capture: the camera or scanner analyzes features such as micro-movements of the face, pupil response, or finger temperature to verify that a real, live user is present. Secondly, the Risk Engine evaluates the metadata associated with the biometric authentication channel. As noted in a Sberbank study, the highest risks are linked to remote authentication scenarios, where the user's device is outside the system's direct control. Accordingly, the system applies differential trust: biometrics captured in controlled environments (e.g., on a corporate device with a certified sensor) are granted a higher level of trust, while those obtained from regular user devices are treated with greater caution. In the latter case, even a successful biometric match may require additional verification to mitigate potential fraud. This approach aligns with NIST recommendations—for example, the SOFA-B document, which assesses the reliability of biometric factors based on the channel through which they are collected.

Thus, the integration of biometrics into adaptive MFA enhances overall system security by adding another barrier for attackers, while incorporating safeguards against biometric-specific threats.

As part of the study, the prototype was tested using a set of experimental scenarios. The evaluation focused on several key metrics: average user authentication time under different conditions, the number of additional verification steps at various risk levels, the number of blocked unauthorized access attempts, and overall user satisfaction. The results confirmed the theoretical assumptions. Under typical conditions (low risk), users logged in with minimal delay: in 85% of cases, access to documents was granted after entering only the primary factor (a password), with no need for additional verification. In contrast, the baseline static MFA system always required extra codes, which increased the average login time by approximately 30%. Thus, in terms of convenience—measured by login time and number of required actions—the adaptive scheme showed a clear advantage.

From a security perspective, simulated attack scenarios demonstrated the system's ability to detect anomalies. Login attempts using clearly stolen passwords from unfamiliar locations were either blocked outright or triggered requests for additional factors inaccessible to the attacker. Separate tests addressed biometric spoofing: attempts to log in using a photo instead of a live face were detected by the system through the absence of liveness indicators, and access was denied, with the event logged as an attack. These results confirm that the combination of the Risk Engine and modern authentication methods can significantly enhance the security of a document-centric system—without compromising convenience for legitimate users.

The proposed approach is applicable to a wide range of systems that require flexible access control to sensitive data—particularly corporate systems, government infrastructures, and cloud-based platforms for exchanging confidential files. In environments where employees work remotely or in hybrid formats, adaptive authentication makes it possible to implement the Zero Trust model in practice, enabling secure document access from anywhere in the world.

Moreover, key components of the system—such as the Risk Engine and the event model—can be integrated into identity and access management (IAM) platforms and broader cybersecurity

solutions. Overall, adaptive multi-factor authentication aligns with the strategy of Continuous Adaptive Trust, which is increasingly recognized as an essential element of modern cybersecurity frameworks. Industry reports indicate that major technology companies—including Google, Amazon, and Microsoft—are already employing risk-based authentication mechanisms to protect user accounts (Wiefling *et al.*, 2021). This demonstrates the practical viability and effectiveness of the concept.

CONCLUSIONS

This paper presents an architecture for adaptive multi-factor authentication in a document-centric system, built on the principles of Zero Trust. The analysis demonstrates that combining the Risk Engine, an event-driven approach, and a diverse set of authentication factors results in a higher level of security compared to traditional models—without placing unnecessary burden on the user. The key strengths of the proposed model lie in its dynamic risk response (adjusting the depth of verification based on contextual factors) and improved user experience for legitimate users (most routine operations in familiar environments require no extra steps). Experimental validation confirmed that the adaptive scheme effectively detects unauthorized access attempts—such as those involving stolen credentials or spoofed biometric data—and substantially reduces the risk of compromising protected documents.

The main findings of the study can be summarized as follows: First, risk-based authentication resolves the long-standing trade-off between security and usability, enabling continuous access control without significantly impairing the user experience (Brodsky, 2018), (Wiefling *et al.*, 2021). Second, incorporating biometric factors into MFA enhances overall protection, but requires the implementation of spoofing countermeasures and robust management of biometric data collection channels (Zakuanova *et al.*, 2018), (Brodsky, 2018). Third, the event model is essential to the effective application of Zero Trust principles: by consolidating and analyzing events from multiple sources—such as authentication systems, applications, and networks—it provides a comprehensive understanding of context, allowing the decision engine to function with greater precision and reliability.

Despite the promising results, the proposed solution has several limitations. First, the effectiveness of the Risk Engine depends heavily on the quality and volume of available data related to user behavior and environmental context. In cases where data is limited—such as with new users or devices—the system may produce false positives (triggering unnecessary MFA prompts) or, conversely, fail to accurately assess risk. Future improvements should focus on refining machine learning models and heuristic risk assessments, as well as collecting more behavioral data to improve reliability over time. Second, introducing adaptability adds complexity to the authentication infrastructure. It requires the integration of multiple components, rule configuration, and compatibility with a range of devices and authentication factors. This increases initial deployment costs and demands skilled personnel to manage and maintain the system. Third, not all organizations or resource types are prepared to adopt the Zero Trust model. For some smaller companies, a traditional static MFA approach may be more practical and cost-effective. As such, adaptive MFA should be viewed as part of a broader organizational security strategy, whose feasibility depends on the risk landscape and available resources.

Future Directions and Standardization – One of the key areas for future work is the development of standardized event models for Zero Trust systems. As previously noted, there is currently no widely accepted specification that defines the format and semantics of authentication and access events within a Zero Trust context (Morrow *et al.*, 2022). Establishing a unified standard would enable

consistent data exchange between the Risk Engines of different solutions, simplify system integration, and improve the reliability of risk assessments through aggregated event data from diverse sources.

Promising directions also include the advancement of adaptive biometric methods, where the system not only validates biometric templates but also dynamically adjusts the required biometric factor or threshold based on contextual factors. Another critical area is the integration of artificial intelligence into the Risk Engine. More sophisticated machine learning algorithms and event correlation techniques will enhance the system's ability to detect complex attacks and anomalies with greater accuracy.

A major objective is to create a self-learning system that incorporates behavioral biometric profiling—analyzing typing dynamics, mouse movement patterns, and individual work styles to continuously refine user identification. To rigorously define system security, it will also be necessary to formalize the models of events and authentication states.

Finally, transitioning from prototype to full-scale deployment will require extensive testing. Evaluating system performance across a broad range of scenarios will enable precise tuning of the balance between security and usability.

REFERENCES

1. Brodsky, A., 2018. Risk-Based Authentication: Balancing Security and Usability (in Russian). *BIS Journal – Information Security of Banks №2(29)/2018*. Available at: https://ib-bank.ru/bisjournal/post/665 [Accessed 10.05.2025].
2. Morrow, T., Popeck, M., & Brown, R., 2022. *Zero Trust Industry Day Experience Paper*. Software Engineering Institute, Carnegie Mellon University. Available at: https://insights.sei.cmu.edu/documents/621/2022_019_001_888817.pdf [Accessed 10.05.2025].
3. NIST, 2020. *SP 800-207: Zero Trust Architecture*. National Institute of Standards and Technology. Available at: https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf [Accessed 10.05.2025].
4. Verizon, 2022. *Data Breach Investigations Report – 2022*. Verizon Enterprise. Available at: https://www.verizon.com/business/resources/T30f/reports/2022-dbir-data-breach-investigations-report.pdf [Accessed 10.05.2025].
5. Wiefling, S., Dürmuth, M., and Lo Iacono, L., 2021 Verify It's You: How Users Perceive Risk-Based Authentication. *IEEE Security & Privacy*, 19(6), pp. 26–36. Available at: https://www.stephanwiefling.de/papers/rba-perceptions-spm2021.pdf [Accessed 10.05.2025].
6. Zakuanova, M.R., Kalinovskii, I.A., 2018. Detection of Spoofing Attacks in Facial Biometric Systems via Texture Analysis (in Russian). In: ITMO University, *Scientific Almanac of ITMO University*, vol. 2, pp. 174–177. Research advisor: Shchemelinin, V.L. Available at: https://science.itmo.ru/wp-content/uploads/2021/08/almanah_2018_tom2.pdf [Accessed 10.05.2025].