IJSES
Volume 9, Issue 10, pp. 5-9, 2025.
International Journal of Scientific Engineering and Science
ISSN (Online): 2456-7361

AI-Driven Adaptive Authentication Framework for Enterprise Systems

Sree Rajya Lakshmi Popury
Senior Engineer Consultant-Systems Engineering, Verizon Communications Inc., Dallas, Texas
Email address: sree.rajya.lakshmi@verizon.com

Abstract— *This paper discusses the adaptive authentication architecture for enterprise solutions, guided by artificial intelligence. The growing threats to static secrets or any other fixed factors, in combination with the high economic damages that can be wrought through their successful exploitation, make it an extremely timely issue. To deliver and prove the architecture of an AI-based adaptive authentication layer, dynamically assessing access risk using user, network, and behavioral metrics, as well as its effectiveness together with scalability within a corporate environment. The novelty of the proposed approach lies in a five-layer architecture that integrates telemetry collection, stream processing via Apache Kafka, an ML-based risk assessment core employing an ensemble of models (gradient boosting, recurrent networks, graph neural networks), a policy engine and an integration layer with IAM/SIEM/SOAR, as well as in the application of federated learning and built-in interpretability mechanisms by BS ISO/IEC 42006:2025. The major results illustrate that this proposed framework delivers not less than eight billion requests per month at sub-second latency with reduced friction by nearly one-third for users and a ninety percent blocking rate of account compromise attempts at false-positive rates of only a fraction of a percent, together with being economically justified because adaptive systems market is projected to grow up to USD 10 billion by 2033. This article will be helpful to information security specialists, enterprise system architects, and developers of authentication solutions.*

Keywords— *Artificial intelligence, adaptive authentication, enterprise systems, Zero Trust, continuous authentication, federated learning, graph neural networks.*

I. Introduction

Static authentication mechanisms rely on unchanging secrets, typically passwords and one-time codes, and thus become primary targets for adversaries. The Verizon DBIR 2024 report notes that credential theft remains the most common initial attacker action, occurring in 24 percent of confirmed security breaches. On average, such credentials appear in one-third of incidents over the past decade [10]. The fallout from leakage or brute-forcing of these secrets is measured not only in data loss: according to IBM's Cost of a Data Breach 2024 report, the global average cost of a single incident reached USD 4.88 million, and stolen or compromised credentials now represent the second most frequent initial intrusion vector [3]. These losses escalate alongside response costs, since each additional minute of access gained via static credentials expands the surface for subsequent damage.

Enterprise environments exacerbate the problem: a distributed workforce, migration of services to the cloud, and widespread API publication increase the attack surface, while internal processes still expect a single, unchanging correct password response without accounting for request context. Even traditional multi-factor authentication raises the bar for attackers only partially, as adversaries increasingly employ push-notification spam, session proxying, and obsolete SMS channels to bypass fixed-factor sets without directly compromising operating systems. The risk continues until verification becomes dynamic and context-dependent. This challenge is what the Zero Trust model addresses, requiring the constant validation of every transaction, regardless of its origin. The private and public sectors are quickly embracing it; in fact, Gartner has predicted that by 2025, more than 60 percent of enterprises will treat Zero Trust as the baseline for their comprehensive security strategy. Still, more than half of them will encounter problems implementing it practically [14]. Therefore, at the core of modern identity architecture, an AI-driven adaptive authentication layer must emerge that links user, device, and network risk assessment to real-time access policies, treating static secrets as only one of many trust variables.

II. Materials and Methodology

The study materials comprised analysis of fifteen sources, including the Verizon DBIR 2024 report [10], the IBM Cost of a Data Breach 2024 report [3], the Gartner Zero Trust forecast [14], market estimates from Viral Visionaries [12], analytical reviews by Viasat [11], MFA statistics from McDade [4], technical white papers by Okta [7], the Okta and Palo Alto Networks case partnership [9], as well as several peer-reviewed articles on Apache Kafka stream processing [1], continuous biometrics [6], federated learning [2, 8], graph neural networks [15] and the BS ISO/IEC 42006:2025 standard [13].

This study therefore takes a multipronged approach, comparing static versus adaptive authentication mechanisms based on metrics such as the prevalence of passwords, adoption of MFA, and contextual factors [4, 5]. A business case for adaptive systems is developed drawing on incident cost data from IBM [3] and market growth projections from Verizon [10] and Viral Visionaries [12]. Technological architectures will be reviewed systematically with performance assessment results for Apache Kafka provided by Alang & Kushwaha, and throughput capability results for Okta's cloud platform available from Okta.

Case study content analysis covered the ML core and graph neural networks for anomaly detection [6, 15]. It also looked at federated learning under GDPR requirements and regional data

processing policies [2, 8]. Normative analysis drew on the BS ISO/IEC 42006:2025 international standard on algorithmic transparency and interpretability [13], as well as integration practices for an adaptive authentication solution with a corporate SIEM/SOAR via the Okta–Palo Alto Networks partnership [9].

Adaptive Authentication Concepts

Adaptive authentication is based on three interrelated concepts: trust levels, authentication factors, and contextual risk. Trust level reflects the system's dynamic confidence that a subject is indeed who they claim to be; it is calculated on a scale from low to high and is recomputed for each access request. Authentication factors are classified as what the user knows, what the user has, or what the user is, and also by the time and place of the request. The latter two are increasingly treated as separate categories, as geocontext and chronocontext allow for the reduction of additional checks without compromising security. According to Figure 1, standard passwords dominate (95 percent), the technology sector leads in MFA adoption (88 percent); administrators (91 percent) substantially outperform ordinary users (66 percent), and among additional factors the most popular are push notifications (29 percent), SMS (17 percent) and soft tokens (14 percent) [4].

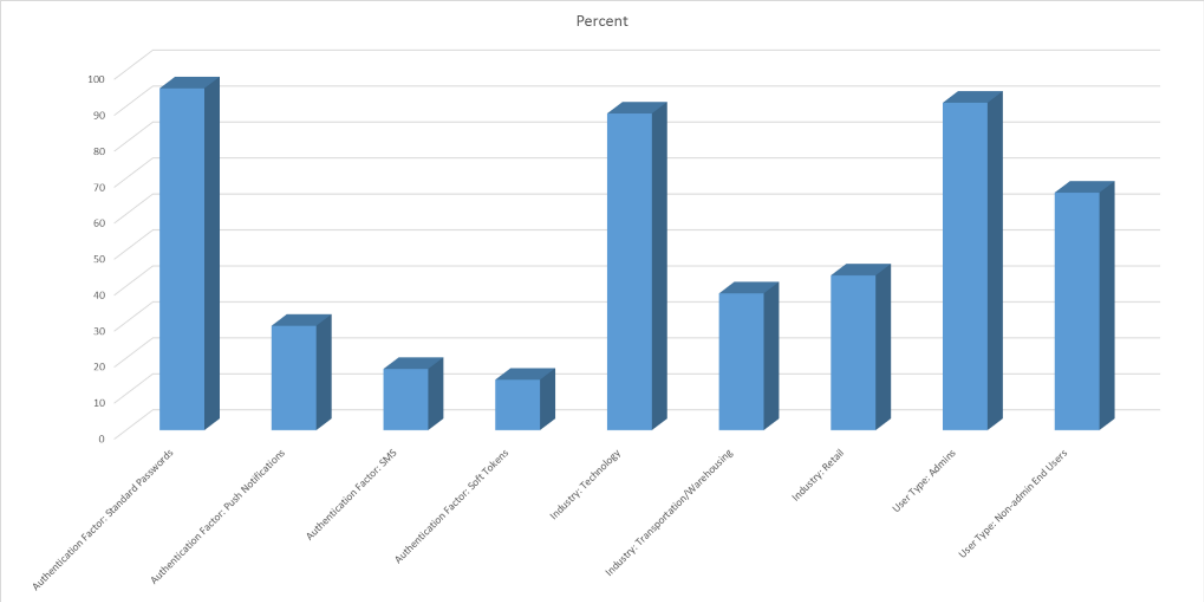

Fig. 1. MFA Adoption Rates by Factor, Industry, and User Role [4]

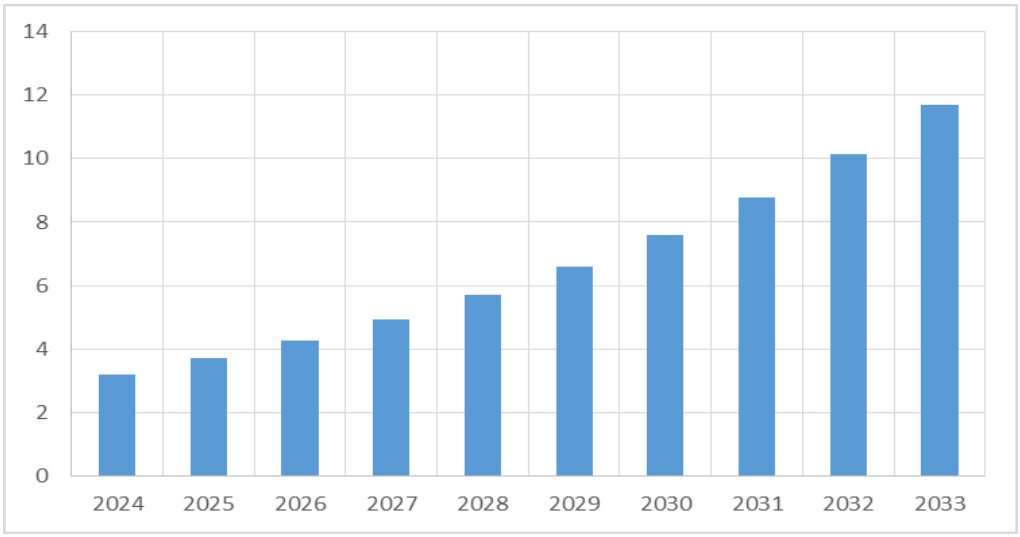

Fig. 2. Global market of adaptive authentication systems [12]

Contextual Risk and Anomaly Response

Contextual risk is defined as the ratio of the detected anomaly to the tolerance for error: if a session originates from a familiar device, risk is almost zero; if, however, login occurs from a new country via a suspicious proxy, risk increases and the system raises the required trust level, activating strong factors such as a FIDO key or biometrics [5].

Economic dynamics confirm the relevance of this approach. The global market for adaptive authentication systems reached USD 3.2 billion in 2024 and, at a sustained annual growth rate of 15.5 percent, will exceed USD 10 billion by 2033, as shown in Figure 2 [12].

This growth is attributed to enterprises' efforts to reduce incident costs without compromising the user experience, and it is precisely the variable degree of verification that provides the optimal balance between security and usability.

The key driver of the adaptive approach is machine learning. Modern implementations collect telemetry on login time, device characteristics, network environment, and micro-behavioral patterns, after which ensembles of algorithms, such as gradient boosting for device scoring and recurrent networks for behavioral biometrics, produce a numerical risk score with virtually no delay. Field practice shows that such systems block up to 90 percent of account takeover attempts in the financial sector, while maintaining false-positive rates at single-digit percentages. Additionally, studies have proven the efficiency of federated learning: model parameters are shared among the branches, and no raw data ever leaves the premises. This changes authentication from a snapshot in time to a form of continuous statistical monitoring that automatically adjusts to legitimate changes in user behavior, as well as threat drift.

In the traditional multi-factor approach, where factors are fixed and applied to every session, user friction is reduced by nearly one-third with adaptive schemes compared to the conventional approach, which provides the same level of protection. However, the principal qualitative shift occurs with the transition to continuous authentication. In this paradigm, behavioral and network features are checked continuously within the session, and access tokens can be revoked immediately upon a rise in risk. Research on Zero-Trust-oriented continuous schemes indicates that the average time to detect a compromise decreases from hours to minutes. In contrast, traditional MFA remains blind after the initial login and is susceptible to push-fatigue attacks. Therefore, the adaptive layer serves as a necessary bridge, rendering multi-factor verification contextual and then smoothly extending it to full continuous authentication, which will be the subject of the subsequent framework architecture analysis.

AI-Driven Framework Architecture

The architecture of AI-driven adaptive authentication comprises five logical layers: telemetry collection, stream data processing, an ML-based risk-assessment core, a policy engine, and an integration layer with enterprise access management services. It is designed for internet-scale loads, as Okta's cloud platform is built to handle. It automatically blocks approximately eight billion attacks per month, thereby establishing the minimal throughput benchmark for any industrial solution [7].

Telemetry begins on the client side. An agent on the device collects a fingerprint of the hardware and software environment. The browser container extracts behavioral signatures, and the network gateway adds information on proxies, VPNs, and IP address reputation. Together, this generates 5 to 10 kilobytes of metadata per access request. In January 2025, the Enhanced Dynamic Zones mechanism in Okta blocked 782 million login attempts originating from residential proxies and concealed VPNs, thereby confirming that even geocontext can serve as a strong risk indicator [7].

The data layer converts the unstructured event stream into a sorted sequence suitable for machine analysis with near-zero latency. In practice, proprietary buses are built around Apache Kafka. Research shows that Kafka's distributed log delivers high throughput with sub-second latency and resilience to individual node failures, remaining a core component for real-time computations in financial, telecom, and IoT scenarios [1]. Such a stream is ingested simultaneously into historical storage and the ML core, enabling a combination of offline training and online inference.

The ML core aggregates incoming features, normalizes them, and feeds them into an ensemble of models: gradient boosting is tuned to static device attributes, recurrent networks handle behavioral dynamics, and graph neural networks detect anomalous relationships among accounts and network segments. Each request is assigned a numerical risk score along with an explainable feature stack. The policy engine translates the risk score into specific actions: at a low level, the request is granted without additional friction; at a medium level, a light WebAuthn check is invoked; and at a high level, a complete factor set is required, or the request is denied outright. The engine also publishes alerts to SOAR or SIEM platforms to complete the response cycle.

The integration layer links the adaptive core to existing IAM systems, proxies, and network-segmentation controls. Typical integration leverages OIDC or SAML standards, allowing external applications to receive access decisions without requiring code changes. In enterprise deployments, it is helpful to automate network responses: the Okta–Palo Alto Networks partnership demonstrates how an identification risk score is automatically forwarded to Cortex XSIAM, triggering session revocation or endpoint quarantine within seconds and thereby implementing the Zero Trust mandate to verify continuously, trust nothing [9]. Such a closed-loop architecture reduces attacker dwell time to minutes and embeds adaptive authentication within the broader digital defense framework.

The framework's algorithmic layer bridges raw telemetry and access control, imposing higher requirements for reliability and speed than traditional authentication mechanisms. The metadata stream—comprising input, session, and network features—feeds into the ML models, where each feature is normalized, aggregated, and converted into a numerical risk estimate in fractions of a second. The triggering threshold is calibrated to ensure that the overall response time remains within acceptable bounds; exceeding these bounds would result in an increased number of user-cancelled operations. This constraint dictates the use of lightweight yet expressive algorithms, notably ensembles of gradient boosting and recurrent networks.

Behavioral biometrics provide the model's primary informativeness: clusters of input dynamics and navigational gestures yield a stable fingerprint that is difficult to imitate. Verification of a sample of continuous sessions showed that for keystroke dynamics, the average False Acceptance Rate does

not exceed 0.5 percent. The False Rejection Rate remains below 1.5 percent at a feature density of approximately twenty per second of interaction, which allowed a reduction in additional factors by nearly one-third without incident growth, as shown in Fig. 3 [6].

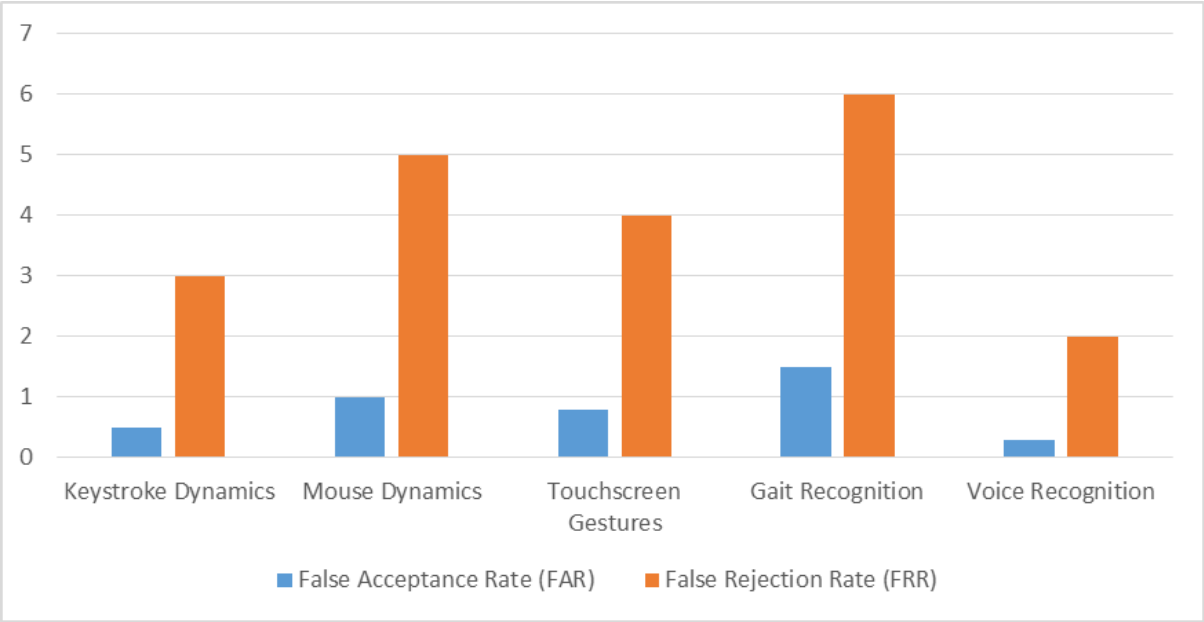

Fig. 3. Comparative Evaluation of False Acceptance and False Rejection Rates in Behavioral Biometric Modalities [6]

The relationships among accounts, devices, and network segments are visualized as a graph, after which graph neural networks uncover topologies typical of bots or compromised users. Recent work presented at IEEE CAI 2025 demonstrated that incorporating a node-controllability metric into edge weights enhances sensitivity to rare anomalies by more than ten points in the F1 metric compared to six baseline methods while maintaining inference time within forty milliseconds, which meets online authentication requirements [15]. This can help stop lateral movement attacks early enough in an industrial environment, as classical rule-based systems will most likely not be able to respond promptly. Federated learning, therefore, comes into play so that the risk-feature repository does not become another cluster of sensitive data. The local clients train fragments of the model directly on their devices. Then only the gradients are sent—thereby bypassing any restrictions on personal data transfers—and also eliminating a bottleneck at a centralized server. It increased the true-positive suspicious login detection rate by five points over a centralized scheme in open F-RBA tests as of December 2024 and reduced the inter-node traffic volume [2]. Other cybersecurity research efforts have also described how distributed training consistently led the way in terms of compliance with regional data laws and internal data-transfer policies [8].

Regulatory Standards and Transparency

Normative pressure for solution transparency is rising in parallel. In July 2025, BSI published the international standard BS ISO/IEC 42006:2025, which requires that every authentication-system inference be accompanied by a verifiable explanation of key features and an audit-context code; this initiative aims to tame the wild west of uncoordinated AI audits and has already gained support from financial-sector regulators [13]. Consequently, the framework's algorithms must include built-in interpretability mechanisms, such as SHAP-style importance estimates, which are logged with each decision and made available for regulatory review without revealing the full model parameters. This balance among effectiveness, confidentiality, and audit transparency completes the shift from static verification to data-driven adaptive authentication.

III. Conclusion

This article demonstrates that the transition from static authentication mechanisms to an AI-driven adaptive approach in enterprise systems addresses both modern cybersecurity challenges and economic efficiency requirements. The analysis shows that credential theft remains a primary attack vector and that incident costs continue to rise, underscoring the need for a more flexible, context-aware approach that goes beyond traditional passwords and fixed MFA factors. The proposed five-layer framework—comprising telemetry collection, stream processing, an ML-based risk-assessment core, a policy engine, and an integration layer—enables adaptive authentication at internet scale. Experimental results demonstrate high throughput for Apache Kafka-based pipelines, low latency for gradient boosting and recurrent network models, and effectiveness of anomaly detection by graph neural networks, such that threats can be detected in milliseconds.

Adaptive systems can easily find their economic validation in market figures: with a size of USD 3.2 billion in 2024, growing at a healthy rate of 15.5 percent per annum, and surpassing the USD 10 billion mark by 2033. One-third of user friction is removed at HIGH, delivering protection while reducing compromise-detection time from hours to minutes,

thereby achieving an optimal balance between security and usability. Future directions federated learning with interpretability embedded for GDPR to insure requirements of new international standards, e.g., BS ISO/IEC 42006:2025 imposing a requisite equilibrium between model efficacy, data confidentiality, and audit transparency particularly financial and governmental domains. Thus, an AI-driven adaptive authentication layer serves as a necessary bridge between traditional multi-factor schemes and full continuous authentication, delivering low false-positive rates alongside high attack-detection efficacy. Further research and industrial deployments will refine optimal model parameters and integration scenarios, fostering the evolution of trust relationships within digital ecosystems.

REFERENCES

[1] K. S. Alang and A. S. Kushwaha, "Stream Processing with Apache Kafka: Real-Time Data Pipelines," International Journal of Research in Modern Engineering & Emerging Technology, vol. 13, no. 3, pp. 216–227, Mar. 2025, doi: https://doi.org/10.63345/ijrmeet.org.v13.i3.13.

[2] H. Fereidouni, H. Senhaji Abdelhakim, D. Makrakis, and Y. Baseri, "F-RBA: A Federated Learning-based Framework for Risk-based Authentication," Arxiv, Dec. 2024, doi: https://doi.org/10.48550/arxiv.2412.12324.

[3] M. Kosinski, "What Is a Data Breach?," IBM, May 24, 2024. https://www.ibm.com/think/topics/data-breach (accessed Jun. 24, 2025).

[4] M. McDade, "A compilation of the latest and relevant statistics regarding multi-factor authentication, its importance, and the wider threat landscape.," Expert Insights, Jan. 10, 2024. https://expertinsights.com/user-auth/multi-factor-authentication-statistics (accessed Jun. 25, 2025).

[5] B. Müller, "Authentication," in Data Protection and Encryption Technology, 2023, pp. 171–176. doi: https://doi.org/10.1007/978-3-031-33386-6_29.

[6] S. Oduri, "Continuous Authentication and Behavioral Biometrics: Enhancing Cybersecurity in the Digital Era," IJIRSET, 2024, Accessed: Jun. 27, 2025. [Online]. Available: https://www.ijirset.com/upload/2024/july/140_Continuous.pdf

[7] Okta, "Okta Secure Identity Commitment," Okta, 2025. Accessed: Jun. 28, 2025. [Online]. Available: https://www.okta.com/sites/default/files/2025-05/okta-secure-identity-commitment-whitepaper.pdf

[8] E. M. Timofte, M. Dimian, A. Graur, D. Balan, and M. Puşcaşu, "Federated Learning for Cybersecurity: A Privacy-Preserving Approach," Applied Sciences, vol. 15, no. 12, pp. 6878–6878, Jun. 2025, doi: https://doi.org/10.3390/app15126878.

[9] D. Todd, "Okta and Palo Alto Networks are teaming up to 'fight AI with AI,'" IT Pro, 2025. https://www.itpro.com/security/okta-and-palo-alto-networks-are-teaming-up-to-fight-ai-with-ai (accessed Jun. 29, 2025).

[10] Verizon, "2024 Data Breach Investigations Report," Verizon Enterprise Solutions, 2024. https://www.verizon.com/business/resources/reports/2024-dbir-data-breach-investigations-report.pdf (accessed Jun. 30, 2025).

[11] Viasat, "Adaptive Authentication in 2024," Viasat, 2024. https://www.rsinc.com/adaptive-authentication.php (accessed Jul. 01, 2025).

[12] Viral Visionaries, "Adaptive Authentication Suite Market Size, Trends, Key Highlights & Investment Outlook," LinkedIn, 2025. https://www.linkedin.com/pulse/adaptive-authentication-suite-market-size-trends-key-4jboc (accessed Jul. 03, 2025).

[13] M. Ward-Brennan, "British body tames AI audit frontier with world's first global standard," City AM, 2025. https://www.cityam.com/british-body-tames-ai-audit-frontier-with-worlds-first-global-standard/ (accessed Jul. 05, 2025).

[14] J. Watts, "Zero trust is moving from hype to reality," Cybersecurity Dive, 2023. https://www.cybersecuritydive.com/news/zero-trust-cybersecurity-Gartner/642399/ (accessed Jul. 07, 2025).

[15] Y. Wei, A. Said, W. Abbas, and X. Koutsoukos, "Robust Anomaly Detection with Graph Neural Networks using Controllability," Arxiv, 2025. https://arxiv.org/abs/2507.13954 (accessed Jul. 09, 2025).