Austrian Journal of Technical
and Natural Sciences, 2026 No 5–6
ISSN 2310-5607

PREMIER Publishing
ppublishing.org

DOI:10.29013/AJT-26-5.6-71-76

SECURITY AS A FEATURE: THE NEXT COMPETITIVE ADVANTAGE IN SOFTWARE PRODUCTS

Aleksandr Polezhaev [1]

[1] Independent Researcher in Software Engineering & AI Director of Digital Innovation Company: Diamdor, Saint Petersburg, Russia

Cite: *Polezhaev A. (2026). Security as a Feature: The Next Competitive Advantage in Software Products. Austrian Journal of Technical and Natural Sciences 2026, No 5–6.* https://doi.org/10.29013/AJT-26-5.6-71-76

Abstract

Contemporary software products face a structural authentication vulnerability that perimeter-based and point-of-entry credential systems cannot resolve: once a session is established, the identity of the operating party is never re-verified. This paper examines the theoretical and empirical basis of that vulnerability, reviews the inadequacy of extant single-modality behavioral biometric approaches, and contextualizes a continuous, passive, multi-modal behavioral authentication framework as a scientifically grounded response. Drawing on cognitive neuroscience, human-computer interaction theory, and machine learning literature, the paper establishes that behavioral individuality in keyboard and pointer interaction is not a statistical artifact but a direct physiological signal exploitable for continuous identity verification. The proposed framework integrates five behavioral signal channels (keystroke dynamics, mouse trajectory, pointer dwell, scroll rhythm, and application transition timing) into a two-tier stacked ensemble classifier with session-level temporal encoding, per-user threshold calibration, and adversarially robust adaptive profile maintenance. Empirical evaluation across 64 users yielded a False Acceptance Rate of 0.28%, a False Rejection Rate of 1.76%, and a mean identity substitution detection latency of 31 seconds, representing statistically significant improvements over all single-modality baselines. The article situates these results within the broader argument that security, when implemented as an architecturally integrated ambient property, constitutes a genuine source of product differentiation and competitive advantage.

Keywords: *continuous authentication; behavioral biometrics; keystroke dynamics; session integrity; machine learning; insider threat; software security; multi-modal fusion; competitive advantage*

Introduction

The commercial software industry has historically treated security as a compliance obligation rather than a product attribute. Authentication systems built around discrete credential verification events (passwords, hardware tokens, one-time codes) provide no mechanism for monitoring the session once access has been granted, leaving the post-authentication period without any form

of identity continuity verification (Bonneau et al., 2012). This design reflects an implicit but consequential assumption: that the identity of the session initiator remains stable throughout the session. The accumulated body of security incident data collected over two decades contradicts that assumption with sufficient consistency to warrant its rejection as an architectural premise.

Behavioral biometrics research provides a theoretically grounded response to this vulnerability. The observation that individuals exhibit stable, measurable, and individually distinctive patterns in their interaction with computing input devices has been documented since at least the 1980s, spanning keystroke dynamics, mouse movement, touch gestures, and gait analysis (Rybnik et al., 2008; Stylios et al., 2022). The core theoretical premise, that motor execution patterns reflect neurologically and physiologically individualized characteristics stable within subjects and discriminative between subjects, is supported by neuroimaging and motor learning research demonstrating that extensively practiced movement sequences consolidate into individual-specific low-level control circuits whose output characteristics are both stable and person-distinctive (Dhawale et al., 2022).

The persistence of this deployment gap is attributable to three distinct limitations of prior systems. Published single-modality systems exhibit Equal Error Rates that remain too high for operational reliance, with keystroke-only implementations typically yielding EERs of 3 to 7% (Morales et al., 2016; Rybnik et al., 2008). Reliance on population-level fixed thresholds, which fail to account for substantial inter-individual variance in behavioral baselines, produces false rejection rates sufficient to generate organizational resistance to adoption (Stylios et al., 2022). The absence of adaptive profile maintenance further prevents such systems from accommodating natural behavioral drift, specifically the gradual change in typing and pointing characteristics that individuals exhibit over extended periods, resulting in progressive false rejection inflation over the deployment lifetime (Bailey et al., 2014). This paper argues that each of these limitations is addressable through an integrated architectural design, and that the resulting system constitutes not merely a technical improvement but a strategic asset for software products that incorporate it. The subsequent sections present the theoretical foundations of the problem, describe the proposed framework, report experimental results, and discuss implications for software product design and competitive differentiation.

Materials and Methods

The discriminative validity of behavioral biometric signals rests on a well-established physiological basis: inter-individual variation in input interaction patterns is a necessary consequence of how motor skills are acquired and consolidated. This is most clearly demonstrated in the neuroscience of motor sequence learning as it applies to typewriting. Research on the neural circuits underlying extensively practiced motor sequences has shown that repeated execution consolidates movement programs into low-level cortical and subcortical control circuits, including the primary motor cortex and the sensorimotor striatum, producing execution patterns that are faster, smoother, and less cognitively demanding than novel movements (Dhawale et al., 2022). Because each individual's keyboard history is unique in volume, style, and temporal distribution, the resulting low-level circuits differ across individuals in ways that are directly observable as inter-individual differences in keystroke timing distributions.

The proposed framework operates as a background software process requiring no dedicated biometric hardware; commodity input peripherals (keyboard and mouse) constitute the sole sensing apparatus. Six cooperating modules form the system: a Signal Acquisition Layer (SAL), a Feature Extraction Engine (FEE), an Ensemble Classifier (EC), an Adaptive Profile Store (APS), a Confidence Scoring Module (CSM), and a Session Integrity Manager (SIM). The SAL captures five concurrent behavioral signal channels: keystroke timing metadata at microsecond resolution without recording key identifiers, mouse pointer trajectory at a minimum of 100 Hz, pointer dwell and click dynamics, scroll event sequences, and application focus transition timing. Raw event buffers are retained exclusively in volatile memory

with a 60-second rolling retention window; no behavioral event sequences are written to persistent storage, satisfying the data minimization principles of contemporary privacy regulation.

At five-second intervals, the FEE extracts a 114-dimensional feature vector from a 30-second sliding window across all five channels, applying quality screening to reject windows with insufficient data density. Per-feature z-score normalization using the enrolled user's own behavioral statistics ensures that the classifier operates on deviation scores expressed in units of the individual's behavioral standard deviation; this per-user normalization scheme renders threshold values interpretable as multiples of the individual's enrollment distribution spread, in contrast to population-calibrated approaches where threshold values carry no user-specific meaning.

The Ensemble Classifier implements a two-tier stacking architecture. The base classifier layer comprises five modality-specific one-class Support Vector Machines trained on each enrolled user's per-modality sub-vectors, and five modality-specific Long Short-Term Memory networks that encode 60 seconds of temporal behavioral context. The session-level temporal encoding provided by the LSTM components addresses a critical limitation of window-level statistical classifiers: replay attacks and session impersonation attempts do not reproduce the characteristic temporal evolution exhibited by genuine users, including the motor warm-up trajectory at session start and fatigue-related drift during extended sessions (Abuhamad et al., 2020). A gradient-boosted meta-classifier trained on 50,000 population-level base-classifier output vectors, including 5,000 adversarially generated mimicry samples, combines the ten base-classifier outputs into a single Platt-scaled probabilistic confidence score.

Per-user threshold calibration replaces the population-level fixed thresholds characteristic of prior systems (Bailey et al., 2014; Morales et al., 2016). The upper threshold T_upper is set at the 5th percentile of the enrolled user's own in-enrollment composite score distribution, ensuring that at most 5% of authentic sessions fall below the threshold under enrollment-equivalent conditions. The lower threshold T_lower is set at 0.60×T_upper, defining a graduated response zone between heightened monitoring and explicit challenge.

The Session Integrity Manager implements a four-state finite automaton with states designated as Active, Watchful, Challenged, and Suspended, each carrying differentiated enforcement actions. The Watchful state increases monitoring resolution without any user-visible action. The Challenged state presents a brief typed text passage for keystroke-based re-verification. The Suspended state blocks session interactions and writes a structured security event to the system log. A configurable Privilege Reduction Mode substitutes graduated access restriction for full suspension in environments where complete session termination would cause service disruption.

Adaptive profile maintenance applies a Mahalanobis-distance-bounded weighted exponential moving average with a learning rate of 0.05 to update the enrolled behavioral model from fully authenticated sessions. A rolling 10-session drift velocity monitor detects adversarial profile poisoning attempts and triggers automatic model reversion upon threshold exceedance.

Performance evaluation was conducted across 64 users, each contributing 15 authenticated sessions and 5 impostor sessions of 60 minutes duration. Single-modality baselines were evaluated using identical classifier architecture and protocol, with four of the five modalities disabled in each baseline condition. The LSTM pre-training population comprised 320 users across 30 sessions each. The meta-classifier training set included 5,000 adversarially generated impostor samples from trained human mimics who had observed target users' behavioral patterns, ensuring that mimicry resistance was tested under adversarially realistic.

Results

The multi-modal framework achieved a False Acceptance Rate (FAR) of 0.28%, a False Rejection Rate (FRR) of 1.76%, an Equal Error Rate (EER) of 1.02%, and a mean identity substitution detection latency of 31.4 seconds. The reduction in EER from

the keystroke-only baseline (3.50%) to the multi-modal framework (1.02%) amounts to a 70.9% improvement relative to the single-modality baseline value, indicating that the addition of four behavioral channels yields error reduction disproportionate to the number of added modalities. The mimicry attack FAR reduction is more pronounced: from 11.2% for the keystroke-only condition to 0.003% for the full framework. This outcome is consistent with the theoretical prediction derived from the independence of the five behavioral channels: if per-modality mimicry success probability is approximately 0.12 for trained human mimics with prior observation access, the joint probability of satisfying all five SVM acceptance criteria simultaneously is approximately 2.5×10^{-5}. The near-correspondence between this theoretical estimate and the empirically observed rate indicates that the five modalities function as substantially independent discriminative dimensions, producing security gains that exceed those achievable through linear combination.

Detection latency improved from 48.3 seconds for the keystroke-only baseline to 31.4 seconds for the full framework, a reduction of approximately 35%. The session suspension false positive rate of 1.6% represents a substantial improvement over the 4.1% rate of the keystroke-only baseline. At 60-minute session granularity, a false positive rate of 1.6% implies that a user conducting four working sessions per day would experience an incorrect suspension event less than once per week on average, a frequency within operationally acceptable bounds for high-security enterprise environments.

Discussion

The empirical results indicate that the proposed framework substantially reduces the authentication gap that point-of-entry credential systems leave open. Session hijacking bypasses Multi-Factor Authentication entirely because the MFA credential is consumed at session initiation and the active session token carries no intrinsic identity binding (Anderson, 2020; Alzahrani et al., 2022). A mean detection latency of 31.4 seconds limits the window during which an identity substitution event can persist undetected, reducing the practical feasibility of data exfiltration within a single hijacked session in typical organizational computing environments.

A productive distinction exists between security as a compliance artifact and security as an ambient product attribute. Compliance-oriented security imposes costs on both vendor and user without producing measurable positive value for either party. Ambient security integrates protective mechanisms that operate without user awareness or active engagement under ordinary conditions (Bonneau et al., 2012). Behavioral authentication, as implemented in the proposed framework, instantiates this category: the system operates within the background of normal computing activity and initiates security events only when behavioral evidence crosses quantitatively calibrated thresholds (Stylios et al., 2022).

The competitive value of this architecture lies in the product's capacity to provide demonstrably stronger post-authentication security guarantees while simultaneously delivering a less obstructive user experience than token-based or challenge-based continuous authentication alternatives (Alzahrani et al., 2022). In market segments where security posture functions as a purchasing criterion, including financial services, healthcare, legal technology, government software, and enterprise productivity tools, the availability of verifiable continuous authentication metrics represents a quantifiable differentiator that can be communicated directly to procurement decision-makers and auditors (Anderson, 2020).

Conclusion

The authentication vulnerability inherent in point-of-entry credential verification is structural in nature: stronger passwords, additional login factors, and more sophisticated session token mechanisms do not address the post-authentication period to the degree required to counter identity substitution, session hijacking, and insider threat activity. The theoretical and empirical literature reviewed in this paper converges on behavioral biometrics as the only passive, hardware-free mechanism available for continuous post-authentication identity verification, grounded

in physiological and cognitive foundations that make behavioral individuality in keyboard and pointer interaction both discriminatively valid and persistently stable over time.

The multi-modal stacked ensemble framework developed by the author addresses the limitations of prior systems through specific architectural decisions with documented empirical effects, yielding a 1.02% Equal Error Rate, a 0.003% mimicry attack FAR, and a 31.4-second mean detection latency. These results demonstrate that the engineering requirements of operational continuous behavioral authentication can be met within the resource envelope of standard enterprise computing hardware. For software products operating in security-sensitive market segments, integration of this capability represents both a measurable advancement in the protective function delivered to users and a sustainable competitive differentiator verifiable through auditable performance metrics.

References

Abuhamad, M., Abuhmed, T., Mohaisen, D., & Nyang, D. (2020). AUToSen: Deep-learning-based implicit continuous authentication using smartphone sensors. IEEE Internet of Things Journal, – 7(6). – P. 5008–5020. URL: https://doi.org/10.1109/JIOT.2020.2975779

Alzahrani, A., Ghosh, U., & Rawat, D. B. (2022). Continuous user authentication using multimodal biometrics and machine learning for zero-trust security. IEEE Access, – 10. – P. 45044–45062. URL: https://doi.org/10.1109/ACCESS.2022.3169972

Anderson, R. (2020). Security engineering: A guide to building dependable distributed systems (3rd ed.). Wiley.

Bailey, K. O., Okolica, J. S., & Peterson, G. L. (2014). User identification and authentication using multi-modal behavioral biometrics. Computers & Security, – 43. – P. 77–89. URL: https://doi.org/10.1016/j.cose.2014.03.005

Bonneau, J., Herley, C., van Oorschot, P. C., & Stajano, F. (2012). The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. Proceedings of the IEEE Symposium on Security and Privacy, – P. 553–567. URL: https://doi.org/10.1109/SP.2012.44

Dhawale, A. K., Wolff, S. B. E., Ko, R., Bhatt, D. L., & Ölveczky, B. P. (2022). Learning-induced changes in the neural circuits underlying motor sequence execution. Current Opinion in Behavioral Sciences, – 48. – 101220 p. URL: https://doi.org/10.1016/j.cobeha.2022.101220

Glasser, J., & Lindauer, B. (2013). Bridging the gap: A pragmatic approach to generating insider threat data. Proceedings of the IEEE Security and Privacy Workshops, – P. 98–104. URL: https://doi.org/10.1109/SPW.2013.47

MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. Human-Computer Interaction, – 7(1). – P. 91–139. URL: https://doi.org/10.1207/s15327051hci0701_3

Morales, A., Fierrez, J., Tolosana, R., Ortega-Garcia, J., Galbally, J., & Gomez-Barrero, M. (2016). KBOC: Keystroke biometrics ongoing competition. Proceedings of the IEEE International Conference on Identity, Security and Behavior Analysis, 1–6. URL: https://doi.org/10.1109/ISBA.2016.7477228

Nurse, J. R. C., Buckley, O., Legg, P. A., Goldsmith, M., Creese, S., Wright, G. R. T., & Whitty, M. (2014). Understanding insider threat: A framework for characterising attacks. Proceedings of the IEEE Security and Privacy Workshops, – P. 214–228. URL: https://doi.org/10.1109/SPW.2014.38

Rybnik, M., Tabedzki, M., & Saeed, K. (2008). A keystroke dynamics based system for user identification. Proceedings of the 7th Computer Information Systems and Industrial Management Applications, – P. 225–230. URL: https://doi.org/10.1109/CISIM.2008.28

Salthouse, T. A. (1984). Effects of age and skill in typing. Journal of Experimental Psychology: General, – 113(3). – P. 345–371. URL: https://doi.org/10.1037/0096–3445.113.3.345

Stylios, I., Skalkos, A., Kokolakis, S., & Karyda, M. (2022). BioPrivacy: A behavioral biometrics continuous authentication system based on keystroke dynamics and touch gestures. Information & Computer Security, – 30(2). – P. 220–240. URL: https://doi.org/10.1108/ICS-12-2021-0212

submitted 20.05.2026;
accepted for publication 04.06.2026;
published 30.06.2026

Contact: aleksandr000polezhaev@gmail.com