Continuous Adaptive Trust Framework for Enhancing Authentication Using Real-Time User Behaviour Analytics

Tharaka Wijekoon

229408C

Master of Science in Computer Science

Department of Computer Science & Engineering
Faculty of Engineering

University of Moratuwa
Sri Lanka

June 2025

Continuous Adaptive Trust Framework for Enhancing Authentication Using Real-Time User Behaviour Analytics

Tharaka Wijekoon

229408C

Thesis/Dissertation submitted in partial fulfillment of the requirements for the degree Master of Science in Computer Science

Department of Computer Science & Engineering
Faculty of Engineering

University of Moratuwa
Sri Lanka

June 2025

DECLARATION

I declare that this is my own work and this thesis/dissertation does not incorporate without acknowledgement any material previously submitted for a degree or diploma in any other University or Institute of higher learning and to the best of my knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgement is made in the text. I retain the right to use this content in whole or part in future works (such as articles or books).

Signature: Tharaka Wijekoon Date: 01/06/2025

The above candidate has carried out research for the PhD/MPhil/Masters thesis/dissertation under my supervision. I confirm that the declaration made above by the student is true and correct.

Name of Supervisor: Prof. Indika Perera

Signature of the Supervisor: Date:

ABSTRACT

Traditional authentication systems struggle to address the dynamic nature of modern cyber threats, often relying on static rules or historical data that fail to adapt to real-time risks. This research proposes a Framework for Continuous Adaptive Trust (CAT) designed to enhance adaptive authentication by integrating real-time user behavior analytics. The framework dynamically assesses contextual factors—including login time, geolocation, device type, and access patterns—to construct behavioral baselines, detect anomalies through hybrid statistical and machine learning models, and enforce adaptive authentication policies. By leveraging a weighted trust score $Ttotal$ that combines behavioral analytics with multi-factor authentication (MFA) outcomes, the system aims to balance security and usability. Integration with the WSO2 Identity Server demonstrates feasibility for enterprise Identity and Access Management (IAM) systems. This work addresses critical gaps in adaptive authentication by prioritizing real-time adaptability, scalability, and privacy-conscious design, offering a foundation for resilient cybersecurity solutions in evolving threat landscapes.

Keywords: Adaptive authentication, user behavior analytics, continuous trust, risk-based authentication, cybersecurity

Table of Contents

List of Figures

List of Tables

List of Abbreviations

Abbreviation	Description
ACM	Association for Computing Machinery
AI	Artificial Intelligence
CAT	Continuous Adaptive Trust
GDPR	General Data Protection Regulation
IAM	Identity and Access Management
IEEE	Institute of Electrical and Electronic Engineers
IoT	Internet of Things
MFA	Multi-Factor Authentication
ML	Machine Learning
OTP	One-Time Password
RBA	Risk-Based Authentication
UBA	User Behavior Analytics
DL	Deep learning
TPR	True Positive Rate
FPR	False Positive Rate
TNR	True Negative Rate
FNR	False Negative Rate
AUC	Area Under ROC Curve
LSTM	Long Short-Term Memory
CCPA	California Consumer Privacy Act
RAT	Remote Access Trojans
SIEM	Security Information and Event Management

1 INTRODUCTION

1.1 Background

In the digital age, securing sensitive data and systems has become paramount as cyber threats continue to evolve in complexity and sophistication. The rapid proliferation of interconnected devices, cloud-based services, and remote work environments has amplified the attack surface, making traditional security measures increasingly obsolete. Among these measures, static passwords—long the cornerstone of authentication—have proven particularly vulnerable. Passwords, often reused across multiple platforms or composed of easily guessable combinations, are susceptible to brute-force attacks, phishing, and credential-stuffing campaigns. High-profile data breaches, such as the compromise of corporate databases containing millions of user credentials, underscore the inadequacy of relying on static, knowledge-based authentication alone. This vulnerability is exacerbated by the human tendency to prioritize convenience over security, leading to weak password practices that undermine even the most robust encryption protocols. As cybercriminals employ advanced techniques like AI-driven password cracking and social engineering, the limitations of static authentication mechanisms have become impossible to ignore.

In response to these challenges, the cybersecurity community has shifted toward more dynamic and context-aware authentication frameworks. Adaptive authentication has emerged as a promising solution, offering a paradigm shift from one-size-fits-all security to risk-based, context-sensitive access control. Unlike traditional methods that apply uniform security measures regardless of circumstances, adaptive authentication evaluates the risk level of each login attempt in real time by analyzing contextual factors such as geolocation, device characteristics, network environment, and time-of-access patterns. For instance, a user logging in from a recognized device within their usual geographic area during typical working hours may undergo minimal authentication steps, such as a single-factor password check. Conversely, an attempt from an unfamiliar device in a foreign country at an unusual hour would trigger additional verification layers, such as biometric scans or one-time passcodes. This dynamic approach balances security and usability, minimizing friction for legitimate users while erecting barriers for malicious actors. However, as noted by Pramila et al. [1], the effectiveness of adaptive authentication hinges on the accuracy of risk profiling, which requires sophisticated algorithms to synthesize behavioral and contextual data into actionable insights.

The evolution of adaptive authentication is deeply intertwined with advancements in User Behavior Analytics (UBA), a discipline focused on modeling and monitoring user activities to detect anomalies. UBA systems employ machine learning algorithms to establish baseline behavior profiles for individual users, capturing patterns such as login frequencies, application usage habits, and transaction

velocities. By continuously comparing real-time actions against these baselines, UBA can flag deviations—such as sudden access to restricted files or atypical data export volumes—that may indicate compromised accounts or insider threats. When integrated with Identity and Access Management (IAM) systems, UBA transforms static authentication into a living, adaptive process. For example, Preuveneers and Joosen [3] demonstrated in their SmartAuth framework how contextual "fingerprints" (e.g., typing speed, GPS data, and Bluetooth proximity) could authenticate users passively, reducing reliance on explicit authentication steps. This fusion of UBA and adaptive authentication enables what researchers term "continuous adaptive trust," where security is not a binary gatekeeper but a fluid evaluation of risk throughout a user's session [4].

Despite its theoretical promise, the practical implementation of adaptive authentication faces significant hurdles. One major challenge lies in the computational and infrastructural demands of real-time analytics. Processing contextual and behavioral data streams—often from disparate sources like endpoint detection tools, network logs, and biometric sensors—requires robust data pipelines and low-latency decision-making engines. Organizations must also address privacy concerns, as the granular monitoring inherent to UBA raises questions about data ownership and user consent. Furthermore, designing adaptive systems that avoid excessive false positives is critical; overly aggressive security measures may frustrate users, eroding trust in the system. For example, a study by Misbahuddin and Bindumadhava [4] highlighted the difficulty of calibrating risk thresholds in machine learning models, where overly sensitive parameters could lock out legitimate users during benign behavioral shifts, such as working late hours during a project deadline. These challenges underscore the need for adaptive authentication frameworks that are not only technically sound but also ethically and ergonomically considerate.

Another layer of complexity arises from the heterogeneous nature of modern IT ecosystems. Enterprises today operate in hybrid environments spanning on-premises servers, cloud platforms, and third-party SaaS applications, each with distinct authentication protocols and logging standards. Achieving seamless adaptive authentication across these silos demands interoperability standards and API-driven integrations that remain works in progress. The LoginRadius report [2] emphasizes that while cloud-based IAM solutions have made strides in unifying authentication processes, gaps persist in correlating risk signals across decentralized systems. For instance, a user's low-risk profile in a corporate email system might not translate to a connected cloud storage service, creating inconsistencies in risk assessment. Addressing these gaps requires industry-wide collaboration to establish shared frameworks for contextual data exchange, as well as advances in federated learning techniques that enable risk models to train on distributed datasets without compromising privacy.

Looking ahead, the maturation of adaptive authentication will likely depend on advancements in artificial intelligence, particularly in unsupervised learning and anomaly detection. Current UBA systems often rely on supervised machine learning models trained on labeled datasets of normal and malicious activities. However, the dynamic nature of cyber threats—where novel attack vectors emerge daily—calls for self-improving systems capable of identifying zero-day anomalies. Researchers are exploring reinforcement learning approaches where authentication systems iteratively refine their risk models based on feedback from successful and thwarted attacks. Additionally, the integration of explainable AI (XAI) principles is critical to ensure transparency in adaptive authentication decisions, particularly in regulated industries where auditors must validate security protocols. For example, if an adaptive system denies access to a user, administrators need interpretable logs detailing which contextual factors (e.g., "unrecognized IP address" or "atypical mouse movement patterns") contributed to the high-risk designation.

The transition from static passwords to adaptive authentication represents a necessary evolution in cybersecurity, aligning defense mechanisms with the fluidity of modern digital threats. By leveraging contextual intelligence and behavioral analytics, adaptive systems promise to mitigate the shortcomings of traditional methods while accommodating the flexibility demanded by today's users. However, realizing this potential requires overcoming technical, operational, and ethical challenges—from building scalable data infrastructure to ensuring user privacy and trust. As organizations navigate these complexities, the collaboration between cybersecurity experts, data scientists, and policymakers will be pivotal in shaping adaptive authentication into a mainstream, practical reality. The foundational work by researchers in risk-based machine learning models [4], contextual fingerprinting [3], and adaptive frameworks [1] provides a roadmap, but the journey toward ubiquitous, seamless, and resilient authentication has only just begun.

1.2 Research Problem

The primary research problem being addressed with this research is the inadequacy of current adaptive authentication mechanisms in providing real-time, robust security against evolving cyber threats [5]. The current mechanisms often fail to accurately assess risk due to their reliance on a predefined set of static rules and historical data, which do not account for the dynamic nature of user behavior and threat landscapes. This research aims to fill this gap by introducing a framework for achieving continuous adaptive trust leveraging real-time user behavior analytics to enhance the effectiveness of authentication in IAM systems.

1.3 Research Objectives

The objectives of this research are as follows:

- Develop a Continuous Adaptive Trust Framework: Create a framework that continuously monitors and analyzes user behavior to detect anomalies in real-time.
- Enhance Adaptive Authentication Mechanisms: Integrate the realtime analytics framework with adaptive authentication systems to dynamically adjust authentication requirements based on current risk assessments.
- Improve Security and User Experience: Ensure that the enhanced adaptive authentication system not only improves security but also minimizes user friction.
- Evaluate Effectiveness: Conduct comprehensive testing to evaluate the effectiveness of the proposed solution in various real-world scenarios.

1.4 Thesis Structure

This thesis is organized into six chapters that systematically address the development, implementation, and evaluation of the Continuous Adaptive Trust (CAT) framework. The structure is designed to guide the reader through the research journey, from foundational concepts to practical validation.

Chapter 1: Introduction

This chapter establishes the research context, highlighting the limitations of traditional authentication systems and the need for adaptive solutions. It defines the research problem, outlines the objectives, and introduces the proposed framework.

Chapter 2: Literature Review

A comprehensive analysis of adaptive authentication techniques, behavioral analytics, and trust management systems is presented. This chapter synthesizes prior

work, identifies gaps in current methodologies, and justifies the need for a hybrid, real-time approach to trust assessment.

Chapter 3: Proposed Method

The core contribution of the thesis is detailed here, including the architecture of the CAT framework. Key components such as behavioral data collection, hybrid anomaly detection (statistical z-score and deep learning models), dynamic risk assessment, and system integration strategies are explained.

Chapter 4: Implementation

This chapter describes the practical deployment of the framework within the WSO2 Identity Server ecosystem. It covers system architecture, data collection modules, behavioral baseline modeling, performance optimizations, and policy enforcement mechanisms. Technical challenges and solutions are discussed.

Chapter 5: Results and Discussion

The framework's efficacy is evaluated through real-world testing and simulations. Metrics include detection accuracy, user experience impact, latency benchmarks, and limitations. Findings are analyzed to validate the framework's ability to balance security and usability.

Chapter 6: Conclusion

The thesis concludes with a summary of contributions, practical implications for cybersecurity, and recommendations for future research. Areas for improvement, such as scalability enhancements and privacy-preserving techniques, are highlighted.

2 Literature Review

Adaptive authentication has emerged as a pivotal technology in the realm of cybersecurity, offering dynamic and context-aware mechanisms to authenticate users. This approach leverages various modalities, including behavioral analytics, to enhance security by assessing the risk profile of each interaction. This literature review delves into the current state of adaptive authentication, focusing on the integration of user behavior analytics to improve its efficacy. The review is structured to provide a comprehensive understanding of the topic, highlighting significant research findings, challenges, and future directions.

2.1 Background

Historically, authentication methods have evolved from simple passwords to more complex multi-factor authentication (MFA) systems. MFA typically combines multiple authentication factors to establish or prove that the person accessing the system is actually who they claim to be. These factors can be divided into three categories, i.e. Knowledge factor - something the user knows (e.g., a password), Possession factor - something the user has (e.g., a smartphone), and Inheritance factor - something the user is (e.g., biometric data). Adaptive authentication builds on this foundation by incorporating contextual data to make real-time decisions about the required authentication factors.

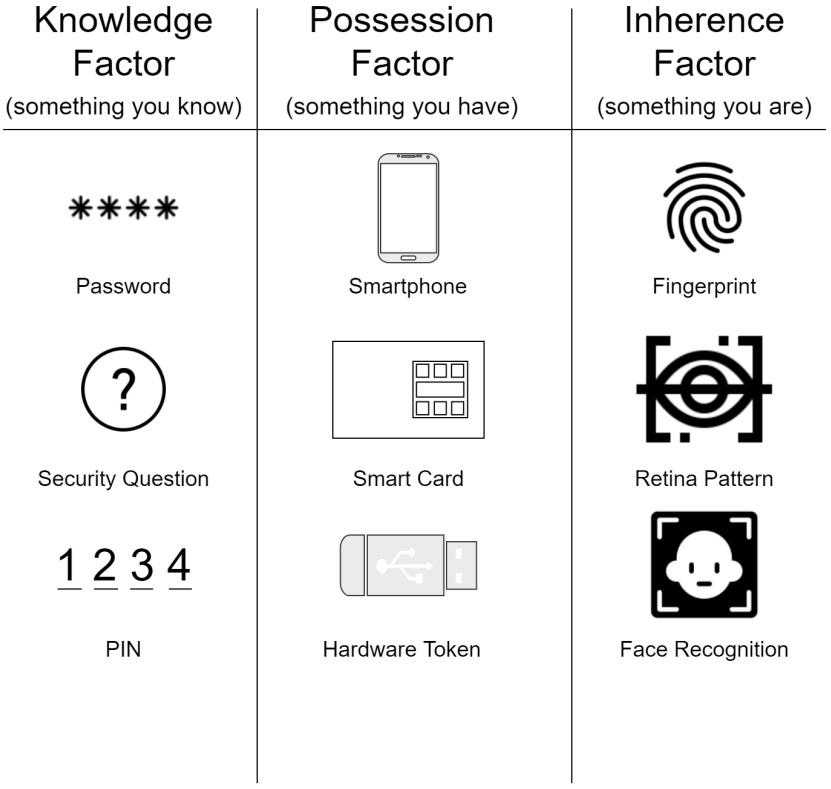

Figure 1: Authentication Factors

Adaptive authentication dynamically selects the best mechanisms among multiple modalities to authenticate a user based on their risk profile, which is generated using behavior and context-based information. This technique is particularly effective in enhancing security for websites and enterprise applications by analyzing large volumes of user, device, and browser data in real-time to generate a risk score that determines the appropriate level of security[1].

Key Concepts in Adaptive Authentication

Contextual Information: This includes data about the user's environment, such as the device being used, the time of access, and the user's location[2]. Context-aware systems utilize various environmental factors—such as location, device proximity, and user behavior—to determine the most appropriate authentication method. For instance, Arias-Cabarcos et al. highlight that adaptive authentication can dynamically select the best mechanisms based on these contextual factors, potentially transforming the password-centric authentication landscape [6]. Similarly, Ryu et al. emphasize the importance of context-aware systems that adapt authentication outcomes based on real-time contextual data, including environmental conditions and user actions [7]. This adaptability not only improves security but also enhances user experience by minimizing unnecessary authentication steps when the context is deemed low-risk.

Risk-Based Authentication (RBA): This involves assessing the risk associated with a user's login attempt based on various factors such as location, device, and behavior[1]. RBA is a specific implementation of adaptive authentication that evaluates the risk associated with a login attempt. Wiefling et al. discuss how RBA can provide robust security with minimal user friction by requiring additional authentication factors only when anomalies are detected in user behavior [8]. This method is particularly effective in large-scale online services, where traditional biometric methods may be impractical due to hardware requirements and user participation [8]. The ability of RBA to monitor user behavior and adapt accordingly makes it a valuable tool in the arsenal of adaptive authentication strategies [9].

Behavioral Biometrics: These are metrics related to human behavior, such as typing patterns, mouse movements, and navigation habits, which are used to verify identity[10]. The integration of machine learning and behavioral biometrics into adaptive authentication systems has shown promise in enhancing security. For instance, touch dynamics and multi-sensor behavior can be leveraged to create implicit authentication systems that continuously verify user identity without explicit user input [10][11]. This continuous authentication approach is particularly relevant in mobile environments, where user interactions can be seamlessly monitored to detect any deviations from established behavioral patterns [12][13]. The use of machine learning algorithms allows these systems to adaptively learn and refine their authentication processes over time, improving both accuracy and security [14].

In addition to these technological advancements, the design and implementation of adaptive authentication systems must consider user perceptions and usability. Research indicates that users may have varying perceptions of security and usability depending on the context in which authentication occurs [15]. Therefore, understanding user behavior and preferences is crucial for developing effective adaptive authentication mechanisms that users are willing to adopt [16].

2.2 Previous Studies and Findings

Several studies have explored various aspects of adaptive authentication, from theoretical frameworks to practical implementations. For instance, Arias-Cabarcos et al. conducted a comprehensive survey on adaptive authentication, highlighting its potential to improve security and usability[17]. Their work emphasizes the need for adaptive systems to consider multiple contextual factors, including user behavior, device information, and environmental conditions.

Another significant contribution is by Abu Bakar and Haron, who discussed the challenges and issues associated with adaptive authentication[18]. They identified several barriers to widespread adoption, such as the complexity of implementation, privacy concerns, and the need for standardized protocols.

The advent of cloud computing has transformed the landscape of identity management, introducing challenges related to trust and security. Bendiab et al. propose a novel blockchain-based trust model aimed at enhancing cloud identity management by addressing security and privacy concerns associated with federated identity management systems [19]. Their work highlights the importance of establishing trust in cloud environments, where poor management can lead to significant security vulnerabilities. Similarly, Ghazizadeh and Cusack emphasize that effective trust management is essential for cloud identity frameworks, as it enables users to make informed decisions regarding security and privacy [20]. This underscores the necessity for continuous adaptive trust mechanisms that can evolve with changing security landscapes.

In the realm of IoT, the dynamic nature of devices and their interactions necessitates an adaptive trust management system. The work of Hamdani et al. illustrates a dynamic distributed trust management scheme tailored for IoT environments, which emphasizes the need for devices to assess trustworthiness based on contextual feedback [21]. This adaptability is crucial for maintaining security in environments where devices frequently join and leave the network. Furthermore, the research by Yanushkevich et al. on cognitive identity management integrates trust assessment with risk evaluation, suggesting that a multi-state dynamical system can enhance trust modeling in IoT contexts [22]. This approach aligns with the need for continuous adaptation in trust management, as it allows for real-time adjustments based on user behavior and environmental changes.

Federated identity management systems also face significant challenges regarding trust. The traditional reliance on centralized identity providers creates vulnerabilities, as highlighted by Shao, who discusses the risks associated with centralized identity management in zero trust networks [23]. This calls for a shift towards decentralized trust models that empower users with greater control over their identity information. Bendiab et al. further explore this by proposing a dynamic trust model for federated identity management that accommodates agile federation establishment [24]. Their findings indicate that continuous adaptive trust mechanisms can facilitate more secure and efficient identity management in collaborative environments.

The integration of context-aware trust management systems is another critical aspect of continuous adaptive trust in IAM. Varadharajan and Nepal propose a context-aware trust management system for IoT applications that operates across multiple domains, emphasizing the importance of contextual information in trust evaluations [25]. This approach aligns with the findings of Awan et al., who advocate for robust distributed trust management mechanisms that account for the dynamic nature of IoT environments [26]. By incorporating contextual factors, these systems can enhance the accuracy of trust assessments, thereby improving security and user experience.

Moreover, the role of risk assessment in trust management cannot be overlooked. Aluvalu et al. propose a risk-aware access control model that integrates trust management into collaborative cloud environments, highlighting the necessity of adapting trust evaluations based on user behavior and risk factors [27]. This perspective is echoed in the work of Hamme et al., who advocate for a hybrid trust model that combines policy and reputation-based approaches to manage distributed trust relationships effectively [28]. Such models are essential for ensuring that trust assessments remain relevant and accurate in the face of evolving threats.

The concept of continuous adaptive trust also extends to the evaluation of trustworthiness in multi-modal authentication systems. The research by Faxin et al. emphasizes the need for dynamic trust assessments that can respond to changes in user behavior and preferences [29]. This adaptability is crucial for maintaining security in environments where users may exhibit varying levels of trustworthiness over time. Additionally, the work of Erat et al. suggests that managerial perceptions and organizational identity play a role in shaping trust dynamics, further complicating the landscape of trust management [30].

The subject of Continuous Adaptive Trust Frameworks utilizing Real-Time User Behavior Analytics encompasses various dimensions of cybersecurity and user trust dynamics. Studies have shown that user trust is significantly influenced by the continuous assessment of behavior and the reliability of systems in maintaining secure interactions.

One pivotal work by Seng et al. discusses an adaptive learning algorithm based on Graph Convolution Networks aimed at understanding implicit trust behaviors among users [31]. Their findings highlight how trust attributes can be inferred from users' interactions, suggesting that adaptive networks can effectively enhance the understanding of user trust dynamics. The use of Jaccard similarity in filtering trust information indicates a methodical approach to improving the quality of trust data utilized in machine learning algorithms, which is foundational in creating adaptive systems capable of real-time responses.

Additionally, the research by Hasegawa et al. outlines cybersecurity interventions within healthcare—a sector where trust is paramount due to the sensitivity of data involved. They emphasize the necessity of an adaptive and responsive approach to cybersecurity, especially in contexts where user behavior may change dynamically due to various external factors [32]. This aligns with findings from Kaur et al., who assert that the implementation of innovative cybersecurity practices, combined with an understanding of trust, is crucial for driving sustainable technology adoption in vulnerable environments [33].

In the realm of mobile commerce, Nguyen and Ha advocate for a deeper understanding of trust and user adaptation, suggesting that trust mediates users' intentions to continue engaging with services, which is fundamental for any adaptive trust framework [34]. By integrating perceived usefulness with adaptive behaviors, the framework allows for a nuanced response to user actions and trust perceptions, enhancing engagement and retention.

Moreover, Singh's investigation into Zero Trust Architecture (ZTA) underscores a paradigm shift towards continuous verification rather than implicit trust [35]. ZTA aligns with the principles of adaptive frameworks, as it necessitates real-time monitoring and reassessment of trust across various interactions. This, coupled with findings from Silva regarding metrics-driven approaches to cybersecurity, further strengthens the relevance of trust adaptations in responding to fast-evolving cybersecurity threats [36]. Silva's work emphasizes the importance of quantitative metrics in assessing trust, which can be translated into algorithms facilitating real-time user behavior analytics.

As user trust continues to evolve in response to changing behaviors and interactions, a dynamic approach that incorporates feedback from user actions—similar to the Bayesian inference approach described by Meng et al.—can improve the detection of potential security threats [37]. This adaptive framework could potentially enhance the reliability of systems while maintaining user confidence, fostering a sustainable environment for digital engagement in sectors heavily reliant on trust.

2.2.1 Current Techniques in Adaptive Authentication

Adaptive authentication techniques aim to balance security and usability by dynamically adjusting authentication requirements based on contextual and behavioral factors. Below is an analysis of prominent approaches, their strengths, and limitations:

1. Static Multi-Factor Authentication (MFA)

Combines two or more authentication factors (e.g., password + SMS OTP) but applies them uniformly across all login attempts.

Strengths:

- Provides higher security than single-factor authentication [1].
- Widely adopted due to simplicity and standardization.

Limitations:

- High user friction: Users face repetitive authentication steps, even in low-risk scenarios [8].
- Inflexibility: Fails to adapt to contextual risks (e.g., trusted devices or locations).

2. Rule-Based Risk-Based Authentication (RBA)

Uses predefined rules (e.g., geo-blocking) to flag risky login attempts and trigger additional authentication steps.

Strengths:

- Reduces false positives in known attack patterns (e.g., logins from blacklisted IPs) [9].
- Lightweight and easy to implement.

Limitations:

- Brittle to novel threats: Rules cannot adapt to evolving attack vectors (e.g., zero-day exploits) [6].
- Limited contextual awareness: Ignores behavioral patterns like navigation habits [5].

3. Behavioral Biometrics

Authenticates users based on unique behavioral traits (e.g., typing rhythm, mouse movements) [10].

Strengths:

- Enables continuous authentication without explicit user input [11].
- Resistant to credential theft (e.g., stolen passwords).

Limitations:

- Privacy concerns: Continuous monitoring raises user distrust [15].

- Computational overhead: Resource-intensive for real-time processing on low-end devices [12].

4. Context-Aware Authentication

Leverages contextual data (e.g., device type, time of access) to adjust authentication requirements [2].

Strengths:

- Reduces friction in trusted contexts (e.g., recurring logins from recognized devices) [7].
- Compatible with IoT and mobile environments [13].

Limitations:

- Static baselines: Relies on historical data, making it vulnerable to sophisticated mimicry attacks [3].
- Limited integration with behavioral analytics for real-time adaptation [6].

5. Machine Learning (ML)-Driven Adaptive Systems

Uses ML models (e.g., neural networks) to analyze user behavior and detect anomalies dynamically [3].

Strengths:

- Adapts to evolving threats through continuous learning [14].
- Detects subtle anomalies (e.g., deviations in navigation paths) [5].

Limitations:

- Data dependency: Requires large, high-quality training datasets [4].
- Complexity: Difficult to interpret model decisions, raising transparency concerns [16].

Table 1: Comparison of Adaptive Authentication Techniques

Technique	Strengths	Limitations
Static MFA	High security, standardized	High user friction, inflexible
Rule-Based RBA	Low false positives for known threats	Brittle to novel attacks, lacks context
Behavioral Biometrics	Continuous authentication, theft-resistant	Privacy concerns, high computational cost
Context-Aware Systems	Reduces friction in trusted contexts	Static baselines, limited real-time adaptation
ML-Driven Systems	Adapts to evolving threats, detects anomalies	Data dependency, complexity

While existing methods improve upon traditional authentication, they suffer from three key gaps:

- Reactive Adaptation: Most systems rely on historical data or static rules, failing to respond to real-time behavioral shifts [6].
- Usability-Security Tradeoff: Techniques like static MFA prioritize security at the cost of user experience, while behavioral biometrics sacrifice transparency for convenience [8].
- Limited Scalability: ML-driven systems and behavioral biometrics struggle with computational demands in resource-constrained environments [12].

These limitations underscore the need for a hybrid framework that combines real-time behavioral analytics, lightweight anomaly detection, and dynamic risk assessment

2.2.2 Challenges and Limitations

The adoption of adaptive authentication, while promising, is fraught with multifaceted challenges and limitations that hinder its widespread implementation and effectiveness. These challenges span technical, ethical, and operational domains, reflecting the complexity of balancing robust security with user convenience and organizational feasibility.

Privacy Concerns

One of the most contentious issues surrounding adaptive authentication is the inherent tension between security and user privacy. The technique relies heavily on the continuous collection and analysis of behavioral and contextual data, such as keystroke dynamics, geolocation, device usage patterns, and network interactions. While this granular data enables dynamic risk assessment, it also raises legitimate concerns about surveillance and data misuse. Users may perceive constant monitoring as intrusive, particularly in jurisdictions with stringent data protection laws like the EU's General Data Protection Regulation (GDPR) or the California Consumer Privacy Act (CCPA). For instance, collecting biometric data (e.g., typing rhythms) without explicit, informed consent could violate privacy norms and legal frameworks. The LoginRadius report [2] emphasizes that organizations must navigate a delicate balance: excessive data collection risks eroding user trust, while insufficient data undermines the system's ability to detect anomalies. Furthermore, the aggregation of behavioral profiles creates lucrative targets for cybercriminals; a breach of such sensitive datasets could expose users to identity theft or targeted social engineering attacks. Mitigating these risks requires robust encryption of stored behavioral data, transparent data governance policies, and mechanisms for users to opt out of specific tracking features—measures that are often resource-intensive and challenging to implement uniformly across global user bases.

Scalability

Scalability remains a critical technical barrier to the deployment of adaptive authentication systems. These systems must process vast streams of real-time data from diverse sources—including endpoint devices, network logs, and cloud services—while maintaining low latency to avoid disrupting user workflows. For large enterprises with millions of users, the computational demands can overwhelm legacy infrastructure, leading to bottlenecks during peak authentication periods. Pramila et al. [1] highlight that adaptive algorithms must not only handle high data volumes but also adapt to fluctuating workloads, such as sudden spikes in login attempts during corporate system updates or global events. Additionally, the heterogeneity of devices and platforms in modern IT ecosystems complicates scalability. For example, behavioral models trained on desktop usage patterns may fail to generalize to mobile or IoT devices, necessitating device-specific adaptations that multiply computational overhead. Cloud-based solutions and edge computing have been proposed to distribute processing tasks, but these approaches introduce their own challenges, such as ensuring consistent security policies across decentralized nodes and managing cross-border data transfer restrictions. Without scalable architectures, adaptive authentication systems risk becoming impractical for large-scale or resource-constrained organizations.

Accuracy

The efficacy of adaptive authentication hinges on the accuracy of its machine learning (ML) models, which must distinguish between legitimate user behavior and malicious activity with minimal error. However, achieving high accuracy is complicated by the dynamic nature of user behavior and the evolving tactics of adversaries. False positives—incorrectly flagging legitimate users as threats—can frustrate users and increase operational costs, as IT teams investigate unnecessary alerts. Conversely, false negatives—failing to detect genuine attacks—compromise security and erode confidence in the system. Addae et al. [38] note that behavioral biometrics, such as mouse movements or app usage patterns, are inherently variable; a user's behavior may change due to stress, fatigue, or new workflows, leading to erroneous risk assessments. For instance, a employee working late hours on a critical project might exhibit atypical login times or rapid data access patterns, triggering unwarranted authentication hurdles. To address this, ML models require continuous retraining with updated datasets that reflect recent behavioral trends, a process demanding significant computational resources and labeled data. Moreover, adversarial machine learning poses a growing threat, where attackers deliberately manipulate their behavior to mimic legitimate users or poison training data. Ensuring accuracy thus becomes an arms race, requiring not only advanced anomaly detection algorithms but also resilience against sophisticated evasion techniques.

Implementation Complexity

The deployment of adaptive authentication systems is a resource-intensive endeavor, often beyond the capabilities of organizations lacking specialized expertise.

Real-time user behavior analytics (UBA) demands integration with existing Identity and Access Management (IAM) frameworks, which may involve overhauling legacy systems not designed for dynamic, data-driven workflows. Hassan et al. [39] underscore the interdisciplinary nature of implementation, requiring collaboration between cybersecurity teams, data engineers, and compliance officers to align technical configurations with organizational policies. Small to mid-sized enterprises, in particular, may struggle with the costs of procuring advanced analytics tools, hiring skilled personnel, and maintaining the infrastructure. Furthermore, the lack of standardized protocols for contextual data exchange complicates interoperability. For example, correlating risk signals from a cloud-based email service with on-premises database access logs may require custom API development, increasing time-to-deployment and potential vulnerabilities. Arias-Cabarcos et al. [40] add that even when technical hurdles are overcome, organizational resistance to change can stall adoption. Employees accustomed to traditional passwords may resist biometric checks or step-up authentication, necessitating extensive training programs and change management strategies. The complexity is further amplified in multinational organizations, where differing regional regulations demand tailored authentication policies, complicating centralized system management.

Synthesis of Challenges

Collectively, these challenges highlight the paradox of adaptive authentication: its strength lies in its contextual granularity, yet this very granularity introduces technical and ethical dilemmas. Privacy concerns and scalability constraints often exist in tension; for example, anonymizing data to protect privacy may reduce the richness of behavioral datasets, impairing model accuracy. Similarly, efforts to simplify implementation—such as using pre-trained ML models—may sacrifice customization, rendering systems less effective for specific user populations. Addressing these limitations requires a holistic approach that prioritizes modular system design, allowing organizations to incrementally adopt features aligned with their risk tolerance and resource capacity. Future advancements in federated learning, edge computing, and explainable AI (XAI) may alleviate some challenges, but until then, organizations must carefully weigh the trade-offs between security, usability, and feasibility when deploying adaptive authentication solutions.

2.2.3 Future Directions

The evolution of adaptive authentication is poised to reshape cybersecurity paradigms, driven by emerging technologies, evolving user expectations, and the relentless sophistication of cyber threats. While current implementations have laid critical groundwork, the future of adaptive authentication hinges on addressing existing gaps while pioneering novel frameworks that harmonize security, usability, and ethical considerations. Four key trajectories—Continuous Adaptive Trust (CAT), standardization, AI/ML integration, and enhanced privacy measures—are central to

advancing this field, each offering transformative potential alongside unique challenges.

Continuous Adaptive Trust (CAT)

Continuous Adaptive Trust (CAT) represents a paradigm shift from static, one-time authentication to a fluid, session-long evaluation of user legitimacy. Unlike traditional models that authenticate users only at login, CAT systems perpetually analyze behavioral and contextual signals—such as navigation patterns, transaction velocities, and device interactions—to dynamically adjust access privileges. For example, a user accessing a financial application might initially undergo multi-factor authentication, but subsequent actions, like transferring large sums to an unfamiliar account, could trigger real-time reauthentication or session termination if deemed high-risk. The LoginRadius report [2] posits that CAT's strength lies in its ability to mitigate "session hijacking" and insider threats, where attackers exploit authenticated sessions to escalate privileges. However, realizing CAT's full potential demands advancements in real-time data processing. Current systems often struggle with latency when correlating disparate data streams (e.g., network logs, biometric sensors), leading to delayed risk assessments. Future CAT frameworks may leverage edge computing to decentralize data analysis, enabling faster decision-making at the device level while reducing reliance on centralized servers. Additionally, CAT must address user experience concerns: overly intrusive monitoring could frustrate users, prompting resistance to adoption. Striking a balance between vigilance and discretion—for instance, by masking risk assessments behind seamless interactions—will be critical to fostering user acceptance.

Standardization

The lack of standardized protocols remains a significant barrier to the interoperability and scalability of adaptive authentication systems. Presently, proprietary solutions dominate the market, creating silos where risk data from one platform (e.g., a cloud service) cannot inform authentication decisions in another (e.g., an on-premises database). Developing universal standards for contextual data exchange, risk scoring, and authentication workflows would enable cross-platform integration, allowing organizations to aggregate risk signals from diverse sources into a unified security posture. Such standards could build upon existing frameworks like FIDO2 (Fast Identity Online) and OAuth, extending them to accommodate behavioral metrics and real-time risk thresholds. Standardization would also democratize access to adaptive authentication, particularly for small and medium enterprises (SMEs) that lack the resources to develop custom solutions. For instance, open-source libraries for risk-based authentication could lower entry barriers, while certification programs might ensure compliance with industry-wide security benchmarks. However, achieving consensus among stakeholders—vendors, regulators, and end-users—will be fraught with challenges. Competing commercial interests and differing regional regulations (e.g., GDPR vs. CCPA) may slow progress, necessitating neutral governing bodies to mediate and promote collaborative innovation.

Integration of AI and ML

Artificial intelligence (AI) and machine learning (ML) are set to revolutionize adaptive authentication by enabling systems to anticipate threats rather than merely react to them. Current ML models excel at detecting known attack patterns but often falter when faced with novel tactics. Future systems could employ unsupervised learning techniques to identify zero-day anomalies, such as subtle deviations in user behavior that evade rule-based detection. For example, reinforcement learning algorithms might iteratively refine risk models based on feedback from thwarted attacks, gradually improving their ability to distinguish between legitimate anomalies (e.g., a user working from a new location) and malicious intent. Addae et al. [38] emphasize the potential of federated learning, where models train on decentralized data without compromising privacy—a critical advantage for multinational organizations handling sensitive regional datasets. However, the reliance on AI/ML introduces vulnerabilities, including adversarial attacks where threat actors manipulate input data to deceive models. Robust defenses, such as anomaly detection ensembles and real-time model validation, will be essential to safeguard these systems. Furthermore, the "black box" nature of advanced ML algorithms poses transparency challenges, particularly in regulated industries requiring auditable decision-making. Explainable AI (XAI) methodologies, which map model decisions to interpretable features (e.g., "access denied due to atypical file download rate"), will be vital to building stakeholder trust and meeting compliance mandates.

Enhanced Privacy Measures

As adaptive authentication systems deepen their reliance on behavioral and contextual data, ensuring user privacy becomes both a technical and ethical imperative. Future research must prioritize privacy-preserving techniques that reconcile data utility with confidentiality. Differential privacy, which injects statistical noise into datasets to prevent re-identification, could anonymize behavioral profiles without eroding their predictive value. Homomorphic encryption, enabling computations on encrypted data, might allow risk assessments without exposing raw user information to third parties. The CyberProtex report [41] highlights zero-knowledge proofs (ZKPs) as another promising avenue, where users can prove authentication claims (e.g., "I am over 18") without revealing underlying data (e.g., birthdates). However, these techniques often entail trade-offs: stringent anonymization may dilute the granularity of behavioral data, reducing detection accuracy. Balancing these trade-offs will require adaptive systems to implement context-aware privacy controls. For instance, a healthcare application handling sensitive patient data might enforce stricter anonymization than a retail platform analyzing shopping habits. Regulatory alignment will also play a pivotal role, as frameworks like GDPR mandate "privacy by design" principles. Future adaptive authentication systems may incorporate dynamic consent mechanisms, allowing users to selectively opt in or out of specific tracking features—a flexibility that could enhance trust without crippling security efficacy.

Converging Pathways

The interplay of these four directions—CAT, standardization, AI/ML, and privacy—will define the next generation of adaptive authentication. For instance, standardized protocols could enable CAT systems to aggregate risk signals from AI-driven, privacy-preserving models across ecosystems, creating a holistic security fabric. However, this convergence demands interdisciplinary collaboration. Cybersecurity experts must partner with data ethicists to ensure AI models respect user autonomy, while policymakers and engineers collaborate on regulatory-compliant standards. The road ahead is not without obstacles: computational bottlenecks, adversarial resilience, and user acceptance loom large. Yet, the foundational work in federated learning [38], behavioral analytics [41], and adaptive frameworks [2] provides a scaffold for innovation. By addressing these challenges through iterative research and cross-sector cooperation, adaptive authentication can evolve from a promising concept into a ubiquitous, resilient pillar of digital security—one that adapts not only to threats but also to the evolving values of a privacy-conscious society.

3 Proposed method

The proposed method for developing a Framework for Continuous Adaptive Trust Using Real-Time User Behavior Analytics comprises several pivotal actions.

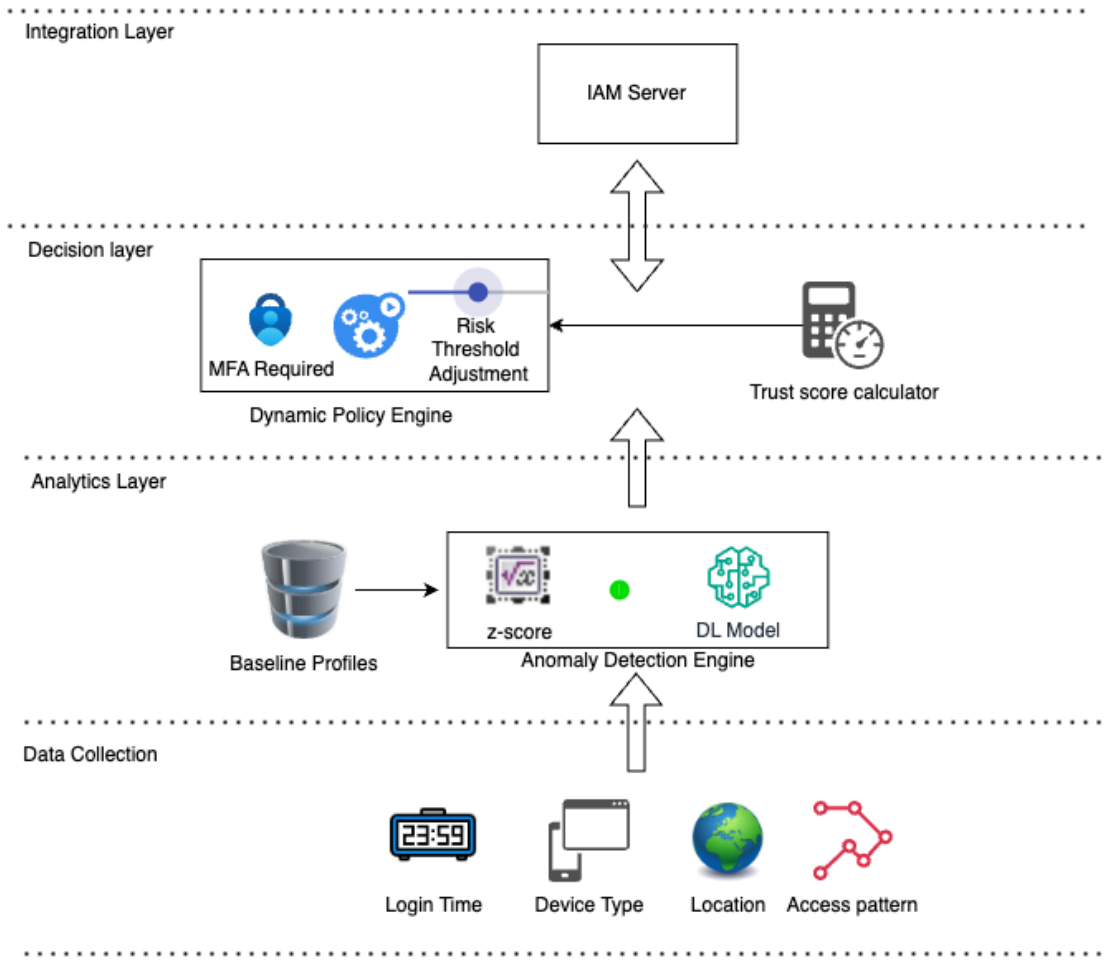

Figure 2: Proposed framework architecture for continuous adaptive trust. The system integrates real-time data collection, behavioral analytics, and dynamic risk assessment with an IAM system.

The method is structured into 5 key phases, each detailed as follows.

3.1 Authentication Data Collection

The foundational step in implementing the proposed framework is the collection of data. This phase involves gathering real-time data on user interactions with the system. With each request to the authentication system the following data attributes can be captured:

- Login Time: Monitor the specific timestamps of user logins to identify common patterns and unusual access times.
- Locations: Collect geolocation data at the time of login, capturing regional and specific location trends, identified through the IP address information.

- Device Type: Identify and record details about the devices used (e.g., desktops, smartphones, tablets), including operating systems and browsers, identified through the user agent.
- Access Patterns: Monitor user behaviors such as typical navigation path, and duration on authentication step.

Table 2: Data Collection Parameters

Feature	Description	Example
Login Time	Timestamp of user login	2025-01-29 14:30:00 UTC
Device Type	Device OS, browser, and hardware details	Android 14, Chrome v120
Location	Geolocation derived from IP address	Colombo, Sri Lanka (6.9271° N)
Access Patterns	Navigation paths, session duration, keystrokes, mouse movements	`/dashboard → /settings` (5min)

Each of the collected data attributes would be a feature, that serves as the foundation for developing user behavior profiles. The features are not limited to the above, more features can be added to improve the framework.

The data collected should be preprocessed to :

- Remove any corrupted, duplicate, or irrelevant records to ensure quality and accuracy in the data.
- Standardize data points, such as login frequency and session duration, to a common scale without distorting differences in the ranges of values.

3.2 Behavioral Analysis and Anomaly Detection

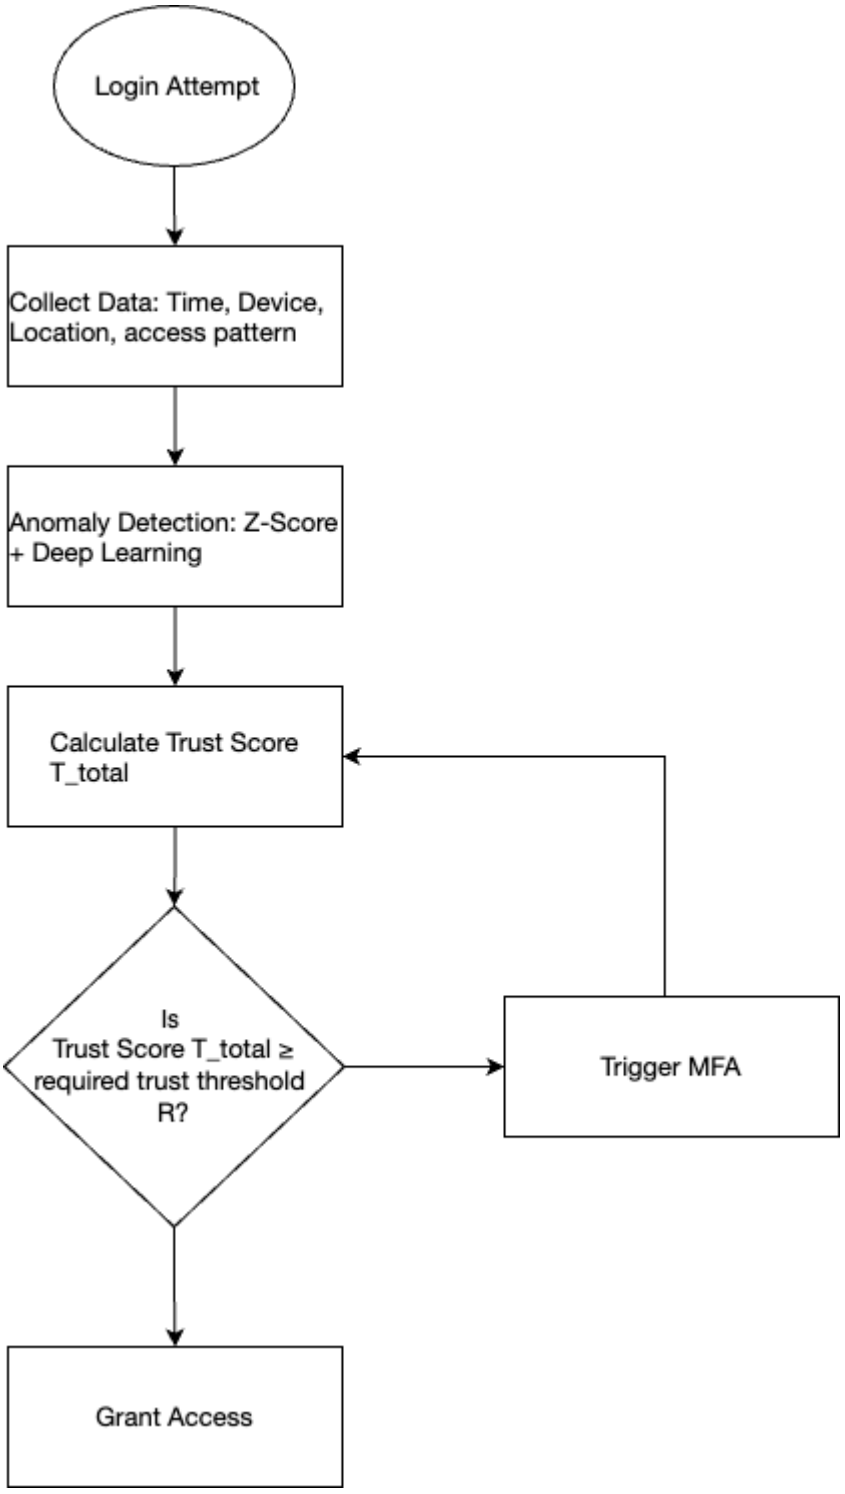

Figure 3: Workflow of trust score calculation. Behavioral data is analyzed against baselines, combined with MFA outcomes, and used to trigger adaptive authentication policies.

In this phase, we construct a baseline model of normal user behavior using the collected behavioral data. The feature vector for the user's interaction at time *t* with the authentication server is:

$Xt = [Login\ time\ t,\ Location\ t\ ,\ Device\ Type\ t\ ,\ Access\ pattern\ t]$

3.2.1 Statistical Anomaly Z-Score

We define the anomaly score Sa(*t*) that measures how much the user's authentication data deviates from the current baseline. To calculate the anomaly score, a z-score approach is used(1),

$$S_a(t) = \frac{|Xt - Xavg|}{Xstd} \quad (1)$$

Where *X*avg is the user's average authentication data, and *X*std is the standard deviation of the user's interactions with the authentication system.

The best approach to combine z-scores for multiple features depends on the feature correlations, goals of the system and computational constraints. The following section evaluates common methods and justify the optimal choice for this framework

1. Maximum Z-Score

$$S_a(t) = \max(Z_1, Z_2, \ldots, Z_n) \quad (2)$$

Strengths:

- Flags the most anomalous feature, useful for systems prioritizing simplicity.

Challenges and Limitations:

- Ignores combined effects of multiple moderate deviations (e.g., Z=1.5 across 4 features would go unnoticed despite a cumulative anomaly).

The hybrid approach (z-score + deep learning) already handles single-feature anomalies via ML. Maximum z-score would underutilize the multi-feature design.

2. Weighted Sum

$$S_a(t) = \sum_{i=1}^{n} w_i Z_i \quad (3)$$

Strengths:

- Allows prioritizing features (e.g., weighting "location" more heavily than "login time" if geolocation is deemed higher risk).
- Flexible and interpretable.

Challenges and Limitation:

- Requires domain expertise to assign weights (wiw i).
- Risks human bias in weight selection.

3. Mahalanobis Distance

$$S_a(t) = \sqrt{(X_t - \mu)^T \Sigma^{-1} (X_t - \mu)} \quad (4)$$

Where Σ is the covariance matrix of features.

Strengths:

- Accounts for feature correlations (e.g., if users logging in late often use tablets, this method adjusts for that relationship).
- Statistically rigorous for multivariate data.

Challenges and Limitation:

- Requires inverting the covariance matrix (Σ^{-1}), which is computationally expensive for real-time systems.
- Suffers from instability if features are highly correlated (e.g., collinearity).

The focus on real-time processing and integration with deep learning makes Mahalanobis less practical due to its computational overhead.

4. Manhattan Norm (Sum of Absolute Z-Scores)

$$S_a(t) = \sum_{i=1}^{n} |Z_i| \qquad (5)$$

Challenges and Limitation:

- Treats all deviations linearly, potentially underweighting extreme outliers (e.g., a z-score of 3 contributes the same as three z-scores of 1).
- Less statistically rigorous for anomaly detection compared to Euclidean norm.

Fails to prioritize high-risk deviations, which is critical for cybersecurity applications.

5. Euclidean Norm (Root Sum of Squares)

$$S_a(t) = \sqrt{\sum_{i=1}^{n} Z_i^2} \qquad (6)$$

Strengths:

- Statistical Rigor: Treats each z-score as a dimension in a multivariate normal distribution. The Euclidean norm represents the Mahalanobis distance if features are uncorrelated and standardized.
- Sensitivity to Outliers: Amplifies large deviations (e.g., a z-score of 3 contributes 9 to the sum, disproportionately flagging extreme anomalies).
- Simplicity: Computationally efficient for real-time systems.

Why we chose it for this framework:

- Euclidean norm is lightweight and aligns with statistical principles for anomaly detection.
- It avoids masking weaker anomalies (unlike the max z-score method) and works well when features are independent or weakly correlated.

If certain features (e.g., "location") are inherently riskier, a weighted sum could complement the Euclidean norm. The Euclidean norm strikes the best balance between simplicity, statistical rigor, and alignment with the hybrid (z-score + ML) architecture. If feature prioritization is critical, a weighted Euclidean norm can be added without significant overhead. This approach ensures the framework remains agile, scalable, and effective in real-world deployments.

3.2.2 Deep learning Model

Z-scores alone may miss subtle or non-linear behavioral anomalies (e.g., gradual deviations in typing speed). To enhance detection capabilities, a deep learning model is employed to identify intricate, non-linear patterns in user behavior. Denote the model's output as $\grave{y}t$, representing the estimated probability that the interaction occurring at time t is genuine.

$$\grave{y}t = f(Xt;\ \theta) \qquad (7)$$

$f(Xt;\ \theta)$ represents the deep neural network model designed to analyze the input features Xt at the given time t.

To train the deep learning model we use the Binary Cross-Entropy loss function:

$$\mathcal{L}(\theta) = -\sum(ytlog\ (\grave{y}t) + (1 - yt)log\ (1 - \grave{y}t)) \qquad (8)$$

yt is the true value which is 1 for legitimate requests, 0 for anomalous requests.

The main advantage of DL is the ability to extract features regarding the problem's requirements, which is a challenge in other ML models. Figure 5 depicts the overall views of traditional ML models and compares them with a DL model.

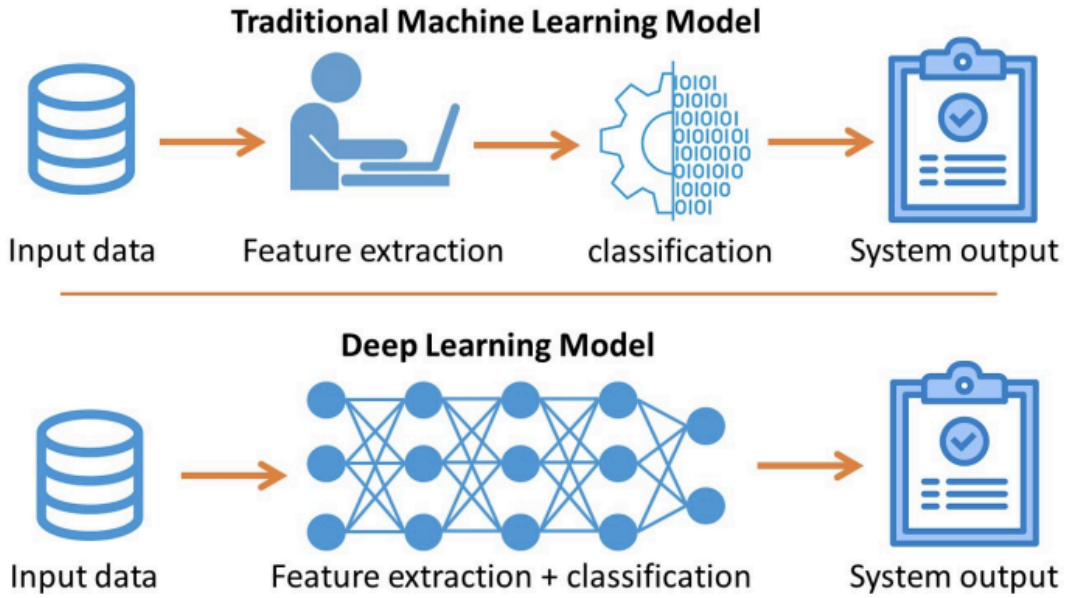

Figure 4: Benefit of using deep learning over traditional ML methods.

A deep neural network (DNN) extends the number of hidden layers in an artificial neural network (ANN) to empower feature extraction ability which is the main advantage of DL over other ML methods. Using multiple layers in DNN can help extraction of higher-level features from the raw input progressively.

3.2.3 Composite Hybrid Anomaly Score

Combine the z-score anomaly (Sa(t)) with the DL model's output ($\grave{y}t$) to create a composite anomaly score.

$$S_{hybrid}(t) = \beta \cdot Sa(t) + (1 - \beta) \cdot (1 - \grave{y}t) \quad (9)$$

$\beta \in [0,1]$ is the weight for z-score vs. DL model.

$1-\grave{y}t$: Converts the DL's "genuine probability" to an anomaly likelihood. Example: If $\grave{y}t$=0.95(95% chance of legitimacy), $1-\grave{y}t$=0.05, indicating low anomaly risk.

We generate an overall trust score combining the results from the behaviour analysis and the MFA results, this trust score is calculated as follows:

The behavioral trust score $Tb(t)$ is inversely related to the anomaly score:

$$Tb(t) = 1 - S_{hybrid}(t) \quad (10)$$

The MFA trust score $T_{MFA}(t)$ is based on the success or failure of the MFA process:

$$T_{MFA}(t) = \begin{cases} 1 & if\ MFA\ succeeds \\ and\ 0 & if\ MFA\ fails \end{cases}$$

The total trust score $T_{total}(t)$ is a weighted combination of these two scores:

$$T_{total}(t) = \alpha Tb(t) + (1 - \alpha) T_{MFA}(t) \quad (11)$$

$\alpha \in [0,1]$ is the weight assigned to behavioral data. α determines the relative importance of the behavioral analysis and MFA results.

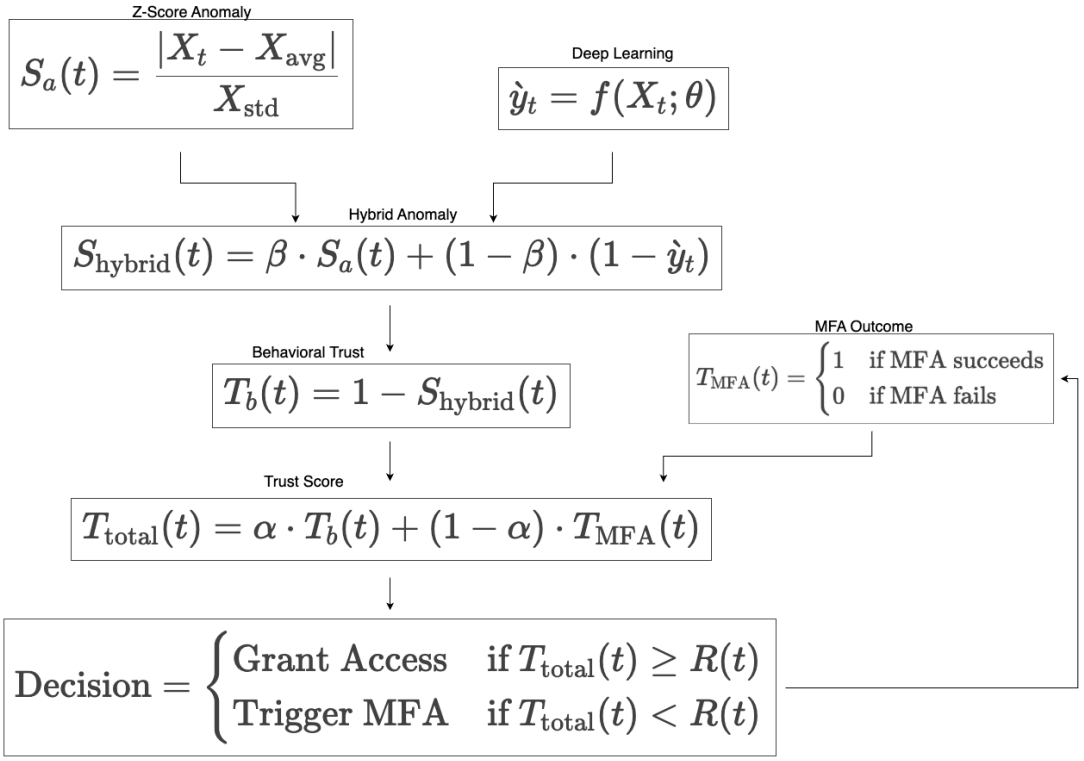

Figure 5: Components of the trust score T_{total} (t), combining behavioral trust (Tb) and MFA outcomes (T_{MFA}).

3.3 Dynamic Risk Assessment

Once the trust score Ttotal (t) is calculated, the system makes an access decision based on the required trust threshold R :

$$A\ decision\ = \begin{cases} Grant\ Access & if\ Ttotal\ (t) \geq R \\ Deny\ Access & if\ Ttotal\ (t) < R \end{cases}$$

For trust scores below the threshold, additional authentication factors will be requested so that the trust score can be improved and the process continues until the users trust score improves and goes beyond the threshold.

With this approach considering the assessed risk level, the system adjusts authentication requirements in real-time. For instance, initiating multi-factor authentication mechanism such as Email OTP if abnormal behavior is detected or reducing authentication burdens to just username and password when normal activity resumes.

This adaptive approach ensures a balance between security and user convenience.

3.4 System Integration

A critical aspect of the proposed framework is its seamless integration with existing Identity and Access Management (IAM) systems.

By implementing and integrating this framework with the open-source WSO2 Identity server[42], we hope to prove the feasibility of this approach and ensure the framework is scalable and compatible with typical IAM systems, allowing for broad applicability and adoption.

This integration enables adaptive authentication processes that are continuously informed by real-time risk assessments.

3.5 Testing and Evaluation

Finally, rigorous testing and evaluation are crucial to assess the efficacy of the proposed framework. This phase includes:

Simulated Testing: Deploy the framework in controlled environments using simulated data sets to evaluate its detection accuracy, responsiveness, and false-positive rates.

Real-World Data Testing: Evaluate performance metrics using actual organizational data, focusing on security enhancement, reduced unauthorized access, and impact on user experience.

Table 3: Planned Evaluation Metrics

Metric	**Description**	**Tool/Method**
Detection Accuracy	% of true anomalies detected	Confusion matrix analysis
False Positive Rate	% of legitimate logins flagged as anomalous	ROC curve evaluation
User Friction	Reduction in MFA prompts for low-risk scenarios	User surveys, session logs
Latency	Time taken for real-time risk assessment	Performance profiling

The outcome of this phase is a validated framework that demonstrates enhanced security capabilities while maintaining efficient user-interaction protocols.

4 Implementation

4.1 System Architecture Overview

The Framework for Continuous Adaptive Trust (CAT) is implemented as an extension of the WSO2 Identity Server, leveraging its native extensibility to avoid the overhead of a standalone backend server. This integrated approach minimizes latency, simplifies deployment, and ensures compatibility with existing authentication workflows. The architecture operates within the WSO2 ecosystem, utilizing adaptive authentication scripts and custom functions to process behavioral data, compute trust scores, and enforce dynamic policies. Figure 6 (Implemented Framework Architecture) illustrates the flow, with all critical components embedded directly into the Identity Server.

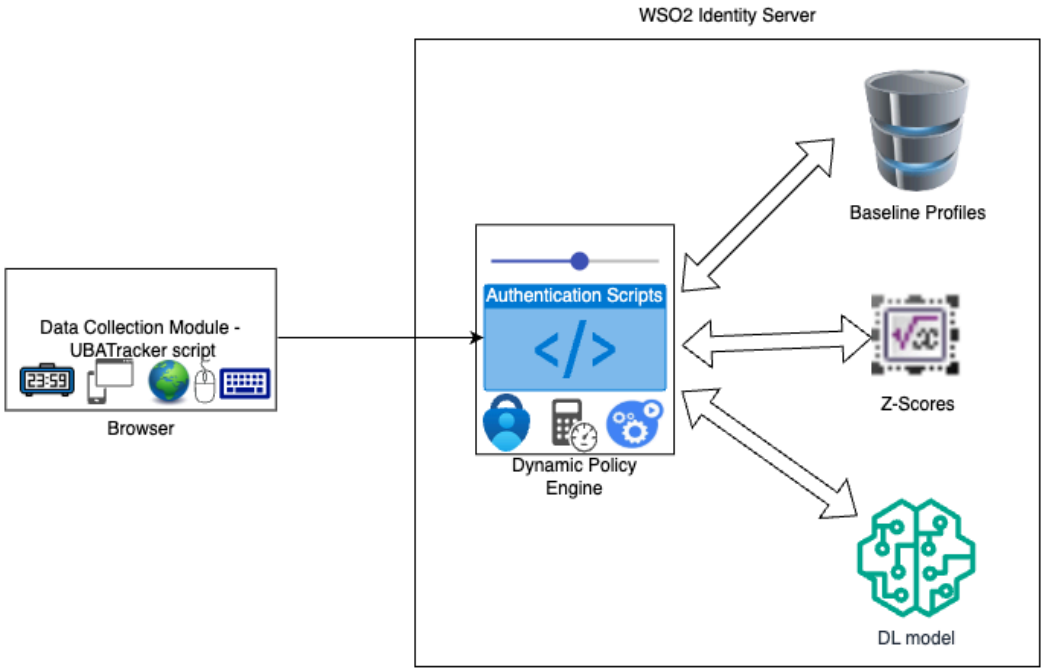

Figure 6: Implemented System Architecture

4.2 Data Collection Module

The Data Collection Module is a critical component of the Continuous Adaptive Trust (CAT) framework, responsible for capturing real-time user behavior data during authentication attempts. This section details the implementation of a JavaScript-based tracking system integrated into the WSO2 Identity Server's login interface (authenticationendpoint/login.jsp).

4.2.1 Implementation Overview

The module employs a client-side script (UBATracker) to collect behavioral and contextual data during user interaction with the login page. The collected data is embedded in the login request as a hidden form field, enabling seamless integration with the WSO2 Identity Server's authentication workflow (Figure 7).

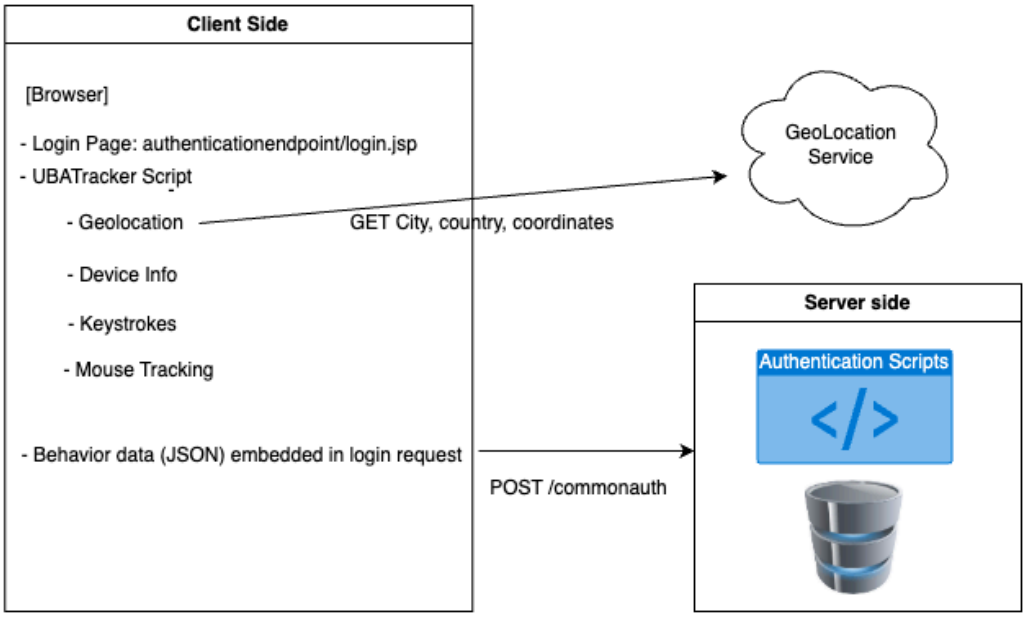

Figure 7: Data flow from client-side collection to server-side processing.

4.2.2 Key Components of the UBATracker Script

1. Initialization

The tracker initializes when the page loads, triggering data collection mechanisms:

```
// Initialize tracker when page loads
document.addEventListener('DOMContentLoaded', () => {
    UBATracker.init();
});
```

- Event Listeners: Track mouse movements and keystrokes.
- Device Profiling: Captures hardware/software characteristics.
- Geolocation: Resolves approximate coordinates via IP address.
- Form Handling: Injects behavior data into login submissions.

2. Behavioral Data Collection

Mouse Movement Tracking

```
document.addEventListener('mousemove', (e) => {
    this.data.mouseMovements.push({
        x: e.clientX,
        y: e.clientY,
        timestamp: Date.now()
    });
});
```

- Captures X/Y coordinates and timestamps.
- Useful for detecting anomalies in navigation patterns.

Keystroke Dynamics

```
document.addEventListener('keydown', (e) => {
    this.data.keyStrokes.push({
        key: e.key,
        timestamp: Date.now()
    });
});
```

- Records typing speed, key press duration, and error patterns.
- Excludes sensitive fields (e.g., password inputs) for privacy compliance.

3. Contextual Data Collection

Device Information

```
captureDeviceInfo() {
  this.data.deviceInfo = {
    userAgent: navigator.userAgent,
    platform: navigator.platform,
    screenResolution: `${window.screen.width}x${window.screen.height}`,
    timezone: Intl.DateTimeFormat().resolvedOptions().timeZone
  };
```

```
}
```

Table 4: Device Information Parameters

Parameter	Example Value	Purpose
userAgent	Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/134.0.0.0 Safari/537.36	Browser Identitfication
platform	MacIntel	OS identification
screenResolution	1728x1117	Device profiling
timezone	Asia/Colombo	Geographic context

Geolocation Resolution

```
async captureLocation() {
    const response = await fetch('https://ipapi.co/json/');
    const locationData = await response.json();
    this.data.location = {
        city: locationData.city,
        country: locationData.country,
        latitude: locationData.latitude,
        longitude: locationData.longitude
    };
}
```

- Uses ipapi.co for IP-to-location mapping.
- Fallback: Server-side geolocation if client-side fails.

4. Data Submission

The collected data is appended to the login form as a hidden field when the login form is submitted, this is done by listening to the submit event of the login form and appending a hidden input - behaviorData to the form. Refer Appendix - UBAScript.js for more details.

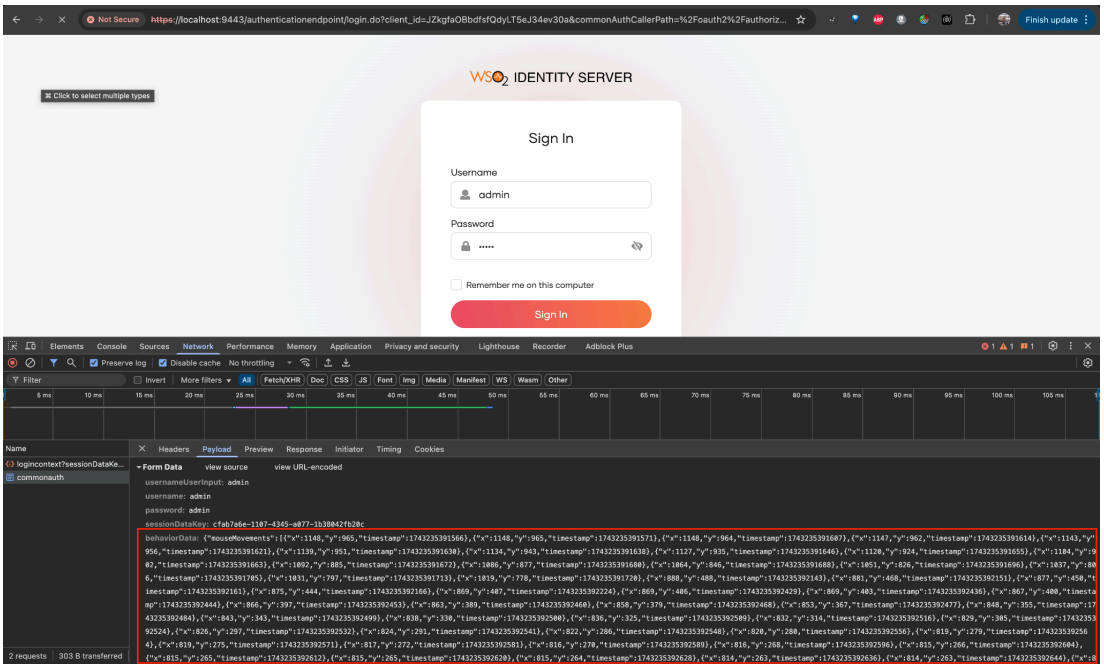

Figure 8: Serialized behavior data embedded in the login request.

4.2.3 Integration with WSO2 Identity Server

The data collection module script is injected into login.jsp to monitor the authentication interface, when the login pages are accessed the script gets initialized and starts collecting the data which is published to the WSO2 server on the login form submissions.

On the server side, WSO2's adaptive authentication scripts extract and process the behaviorData parameter that is passed on as hidden input during the authentication:

var behaviorData = context.request.params.behaviorData[0];

4.3 Behavioral Baseline Modeling

This section details how user behavior data is stored, baseline requirements are enforced, and adaptive sliding window mechanisms enable continuous model refinement.

4.3.1 Baseline Establishment Workflow

The CAT framework remains inactive for users until sufficient behavioral data is collected to establish a reliable baseline. The behavioral data is stored per user in the database, and a user behavior profile is maintained to track the baseline completion status for each user. These tables are newly created in the WSO2 Identity database to support the CAT framework.

1. IDN_USER_BEHAVIOR_DATA: Stores raw behavioral metrics per login attempt.
2. IDN_USER_BEHAVIOR_PROFILE: Tracks baseline completion status and login counts.

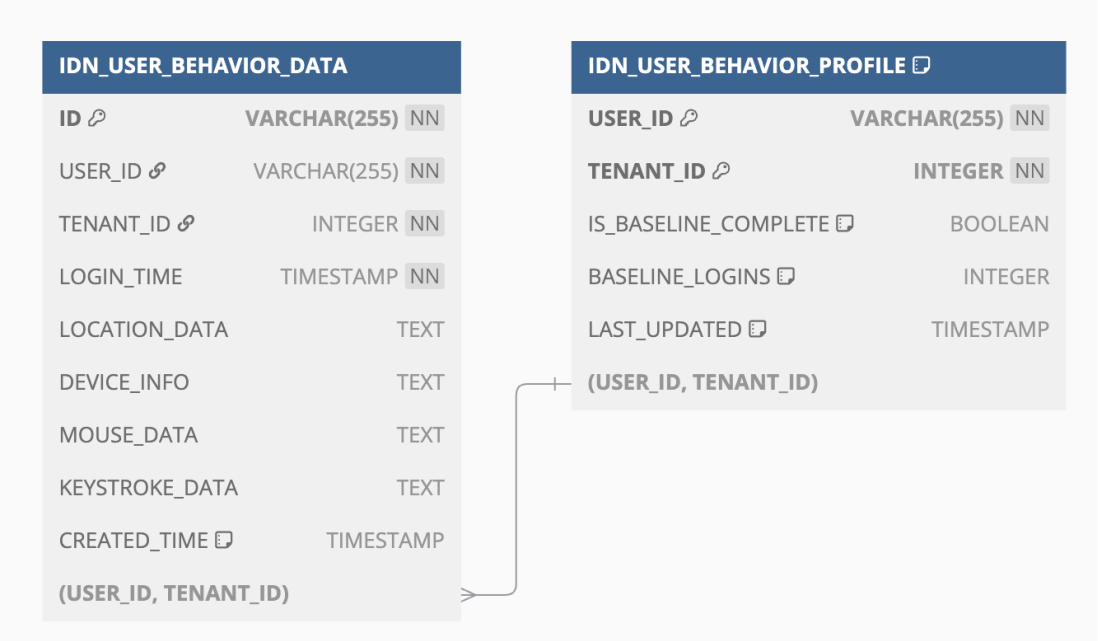

Figure 9: New tables for the behavior baseline tracking

Once enough behavioral data is collected for the user to complete a window to calculate the baseline profile, the framework gets activated addressing the cold start problem. The default window size is defined as 100, but is a configurable value for the CAT framework.

During the baseline establishment stage the user would be prompted with the MFA regardless of whether the login attempt is genuine or not.

4.3.2 Sliding Window Baseline

Once the framework gets activated after collection of behavioral data to fit the window size, the adaptive authentication scripts on the Identity server gain access to the window to establish the baseline for the behavior data.

With each authentication attempt the window gets updated and the adaptive authentication scripts have access to the latest data in the window to establish the baseline. The baseline keeps on updating with the sliding window.

This implementation ensures the CAT framework adapts to gradual behavioral changes while maintaining computational efficiency. The sliding window approach provides inherent protection against concept drift, making the framework resilient to evolving user behavior patterns and sophisticated mimicry attacks.

4.4 Data Processing

4.4.1 Z-Score Standardization

Features like "login time" (0–23), "location" (distance in km), and "keystroke patterns" (milliseconds) are on vastly different scales. Without standardization, one feature (e.g., location) could dominate the anomaly score Sa(t), masking smaller deviations in others (e.g., device type).

Z-scores transform all features into unitless measures of deviation, enabling fair aggregation into Sa(t). Standardization ensures the deep learning model (trained on normalized data) can generalize effectively.

4.4.2 How each feature is standardized

Table 5: Feature standardization

Feature	Type	Standardization Method
Login Time	Continuous numeric	Compute $Z = \frac{\text{Time} - \mu_{\text{Time}}}{\sigma_{\text{Time}}}$
Location	Geospatial distance	$Z = \frac{\text{Distance (km)} - \mu_{\text{Location}}}{\sigma_{\text{Location}}}$
Device Type	Categorical	Use frequency encoding: Convert device type to its historical usage probability (e.g., 0.7 for mobile), then standardize with Z-Scores.
Access Patterns(Mouse + Keystroke)	Time-series/behavioral	Standardize with Z-Scores - velocity, time between keystrokes.

Practical example of how the feature standardization and anomaly scores are calculated for each feature.

Baseline Statistics (Hypothetical User):

Table 6: Baseline statistics for hypothetical user

Feature	Mean (μ)	Std (σ)
Login Time (hour)	9 (9:00 AM)	2 hours
Location (km)	0 (usual location)	50 km
Device Type (mobile)	0.7 (70% usage)	0.15
Mouse Speed (px/s)	500	100
Keystroke Dwell (ms)	200	50

Current Authentication Data (Xt):

Login Time: 13 (1:00 PM)

Location: 150 km away

Device Type: Tablet (historical frequency = 0.1)

Mouse Speed: 700 px/s

Keystroke Dwell Time: 300 ms

Standardized Scores:

Login Time:

$$Z_{\text{time}} = \frac{13-9}{2} = 2.0$$

Location:

$$Z_{\text{location}} = \frac{150-0}{50} = 3.0$$

Device Type:

$$Z_{\text{device}} = \frac{0.1-0.7}{0.15} = 4.0$$

(Negative score indicates rare usage of tablet.)

Mouse Speed:

$$Z_{\text{mouse}} = \frac{700-500}{100} = 2.0$$

Keystroke Dwell:

$$Z_{\text{keystroke}} = \frac{300-200}{50} = 2.0$$

These calculations are carried out using custom adaptive authentication functions written in Java deployed on the WSO2 Identity server.

4.5 Hybrid Anomaly Detection Model

The hybrid anomaly detection model is a combination of the statistical z-score model and the deep learning model.

4.5.1 Statistical Z-Score Model

The foundation of our adaptive trust evaluation system includes a robust statistical model that quantifies behavioral deviations from established user patterns. This implementation employs Z-score standardization, a well-established statistical method for identifying outliers within multivariate behavioral data.

The statistical model plays a dual role in our hybrid approach:

Independent verification: Provides a complementary trust evaluation method based on well-established statistical principles

Training supervision: Supplies ground truth labels during the early training phases of the deep learning model

Fallback mechanism: Ensures reliable trust evaluation when insufficient data exists for deep learning analysis

4.5.1.1 Multi-dimensional Behavioral Analysis

Our implementation analyzes four key behavioral dimensions to construct a comprehensive trust profile:

Location Analysis - The system evaluates the geographical consistency of user login locations using the Haversine formula to compute distances between current and historical coordinates. The Haversine distance calculation accurately accounts for the spherical geometry of Earth.

Login time Analysis - The system analyzes login time patterns, capturing time-of-day variations. Historical time patterns are maintained as frequency distributions, allowing the system to identify unusual login times.

Device Information Analysis - Device consistency is evaluated by comparing browser fingerprints, operating systems, and hardware characteristics.

Access Pattern Analysis - The system examines sequences of keystrokes and mouse velocities to identify any anomalies.

4.5.1.2 Euclidean Norm Score Aggregation

The aggregation of individual dimension scores using the Euclidean Norm (Root Sum of Squares) method is carried out to obtain the combined trust score.

```
// Calculate statistical score using Euclidean Norm (Root Sum of Squares)
double statisticalScore = Math.sqrt(
    Math.pow(locationScore, 2) +
    Math.pow(timeScore, 2) +
    Math.pow(deviceScore, 2) +
    Math.pow(accessPatternScore, 2)
)
```

This approach offers several advantages over simple averaging as it properly accounts for the multivariate nature of the data giving appropriate weight to large deviations in any dimension. It maintains the statistical interpretation of distance in multidimensional space and provides a normalized score between 0 and 1, with higher values indicating greater trust.

4.5.2 Deep Learning Model

The implementation of our deep learning model for user behavior analysis employs a hybrid recurrent neural network architecture tailored specifically for multimodal behavioral data. Unlike traditional feature engineering approaches, our model directly processes raw behavioral sequences, allowing it to learn complex patterns that may not be evident through manual feature extraction.

The architecture consists of three primary components:

1. A sequence processing pathway for mouse movement data
2. A parallel sequence processing pathway for keystroke dynamics
3. A static feature processing component for location and device information

These pathways are integrated through a fusion layer that combines the representations before final classification..

4.5.2.1 Sequential Data Processing

Mouse Movement Analysis

Mouse movements are processed through an LSTM (Long Short-Term Memory) layer to capture temporal dependencies in user interaction patterns. Raw mouse coordinates and timing information are normalized and structured as follows:

```
// Normalization of coordinates to [0, 1] range
mouseData[0][0][i] = x / screenResolution.width;
mouseData[0][1][i] = y / screenResolution.height;

// Time delta normalization, capped at 2 seconds
double timeDelta = Math.min((timestamp - previousTime) / 1000.0, 2.0) / 2.0;
```

This sequential representation preserves trajectory information, velocity patterns, and acceleration profiles that characterize individual user behavior.

Keystroke Dynamics

The keystroke sequence processing branch employs a similar LSTM architecture to analyze typing patterns. Each keystroke is represented by:

- A normalized hash of the key identity
- Inter-key timing intervals

This approach captures rhythmic typing patterns unique to each user:

```
// Hash the key to a normalized value
keyHash = (Math.abs(key.hashCode()) % 1000) / 1000.0;

// Normalize time delta between keystrokes
double timeDelta = Math.min((timestamp - previousTime) / 1000.0, 2.0) / 2.0;
```

4.5.2.2 Static Feature Integration

Complementing the sequential data, our model incorporates static contextual features:

- Geospatial coordinates (normalized latitude and longitude)
- Temporal patterns (hour of day, day of week)
- Device characteristics (user agent, platform, screen resolution)

These features provide essential context for interpreting the dynamic behavioral sequences:

```
// Normalize latitude (-90 to 90) to [0, 1]
staticFeatures[0] = (location.getDouble("latitude") + 90) / 180.0;
// Extract hour of day and normalize
staticFeatures[2] = hour / 24.0;
```

4.5.2.3 Network Architecture and Training

The neural network architecture was implemented using DL4J (DeepLearning4J), configured with:

- LSTM layers with 32 and 16 units for mouse and keystroke processing respectively
- A fusion dense layer with 20 units and ReLU activation
- A binary output layer with sigmoid activation for trust score prediction
- Binary cross-entropy loss function optimization

```
MultiLayerConfiguration conf = new NeuralNetConfiguration.Builder()
    .seed(123)
    .updater(new Adam(0.001))
    .l2(1e-5)
    .weightInit(WeightInit.XAVIER)
    .list()
    .layer(0, new LSTM.Builder()
        .nIn(MOUSE_FEATURES)
        .nOut(32)
        .activation(Activation.TANH)
        .build())
    // Additional layers...
    .build();
```

The model employs adaptive learning rates through the Adam optimizer with L2 regularization to prevent overfitting.

4.5.2.4 User-Specific Model Training

A key innovation in our approach is the per-user model training methodology. Rather than employing a single global model, we maintain separate models for each user, allowing for personalized behavior analysis:

1. Each user's model is trained incrementally as new behavioral data is collected
2. Statistical scores serve as initial ground truth during early training phases
3. A minimum training sample threshold (n=10) ensures reliable predictions

This approach addresses the cold start problem while enabling continuous adaptation to evolving user behavior patterns.

4.5.2.5 Model Persistence and Efficiency

The implementation includes an efficient model persistence mechanism that serializes trained models to a database:

- Models are cached in memory (limited to 100 concurrent models)
- Periodic persistence to database storage after training iterations
- Lazy loading of models when needed for prediction

This approach balances computational efficiency with scalability for multi-tenant environments.

The IDN_USER_BEHAVIOR_MODELS table was created in the Identity server database to store the user models.

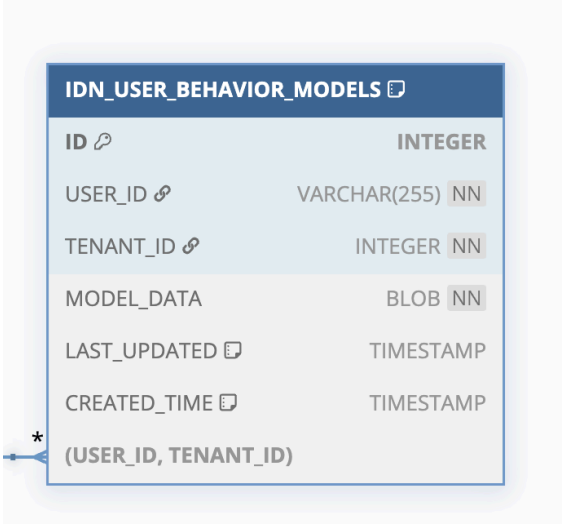

Figure 10: New table created to storing the model data

4.6 Dynamic Policy Enforcement

The Dynamic Policy Enforcement module is the decision-making core of the Continuous Adaptive Trust (CAT) framework, responsible for enforcing Multi-Factor Authentication (MFA) based on real-time trust scores. This section details the implementation of adaptive authentication policies using WSO2 Identity Server's script-based authentication framework.

4.6.1 Policy Enforcement Workflow

The enforcement logic follows a risk-based approach, dynamically escalating authentication requirements when behavioral anomalies are detected. The workflow is governed by two critical parameters:

Trust Threshold (trustThreshold): Minimum acceptable trust score (default: 0.7).

Alpha (α): Weighting factor balancing behavioral vs. MFA scores (default: 0.6).

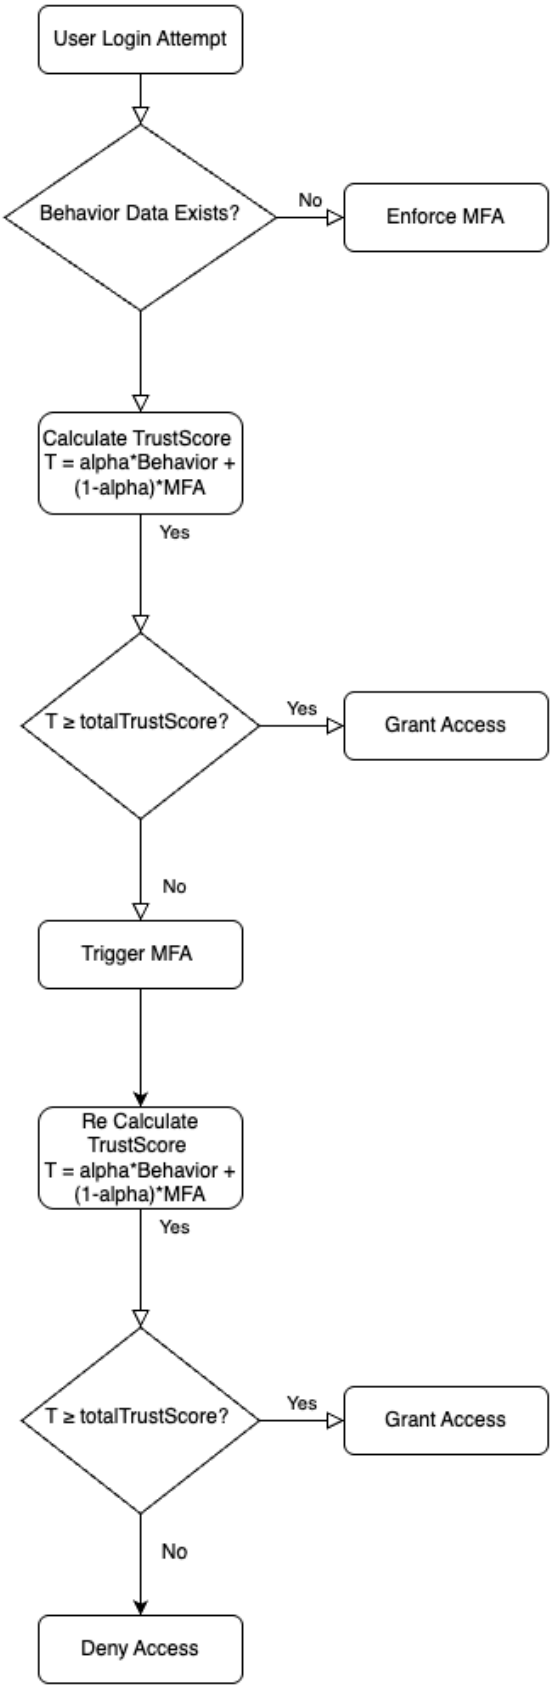

Figure 11: Conditional logic for MFA enforcement based on trust score thresholds.

4.6.2 Authentication Script Structure

The WSO2 Identity server allows defining an authentication script which is written in javascript to define how the authentication can be handled. We use these authentication scripts to dynamically enforce MFA based on the trustscore thresholds.

```
var trustThreshold = 0.7;  // Configurable threshold
var alpha = 0.6;           // Behavioral vs MFA weight

var onLoginRequest = function(context) {
   executeStep(1, { // Step 1: Basic authentication
      onSuccess: function(context) {
         // Main policy logic
      }
   });
};
```

4.6.3 Trust Score Calculation

The overall composite trust score combines behavioral and MFA components, since initially for the first step of the authentication process the MFA is not prompted an MFA score of zero is given.

```
var behavioralScore = analyzeBehavior(behaviorData, historicalData); // [0-1]
var mfaScore = 0; // Initial state

// Initial trust calculation
var totalTrustScore = (alpha * behavioralScore) + ((1 - alpha) * mfaScore);
```

The behavioralScore is the trustscore that is derived from combining the statistical and deep learning models, it is combined with the binary MFA score(0=not completed, 1=verified) to determine the total trustscore.

4.6.4 Conditional MFA Enforcement

```
if (totalTrustScore >= trustThreshold) {
   Log.info("Allow basic authentication");
} else {
   executeStep(2, { // Trigger MFA
      onSuccess: function(context) {
         mfaScore = 1;
         // Recalculate with MFA verification
         totalTrustScore = (alpha * behavioralScore) + ((1 - alpha) * mfaScore);

         if (totalTrustScore < trustThreshold) {
            sendError('High risk login'); // Final denial
         }
```

```
    }
});
```

After MFA verification the mfaScore gets set to 1 and the totalTrustScore is recalculated, This gives MFA outcomes 40% weight in final decision if the default 0.6 alpha is used. If the totalTrustScore does not go up beyond the threshold even with the MFA, the login is marked as a high risk login and denied with an error.

Administrators can customize enforcement through tuning the following configurations in the script.

trustThreshold: The higher the trust threshold, the better the stricter the policy enforcement would be.

alpha: Increasing the alpha value would priority the behavioral analytics score over the MFA score.

Step Chaining: More MFA steps could be introduced to the authentication follow, this would allow improving the totalTrustScore in high risk login scenarios.

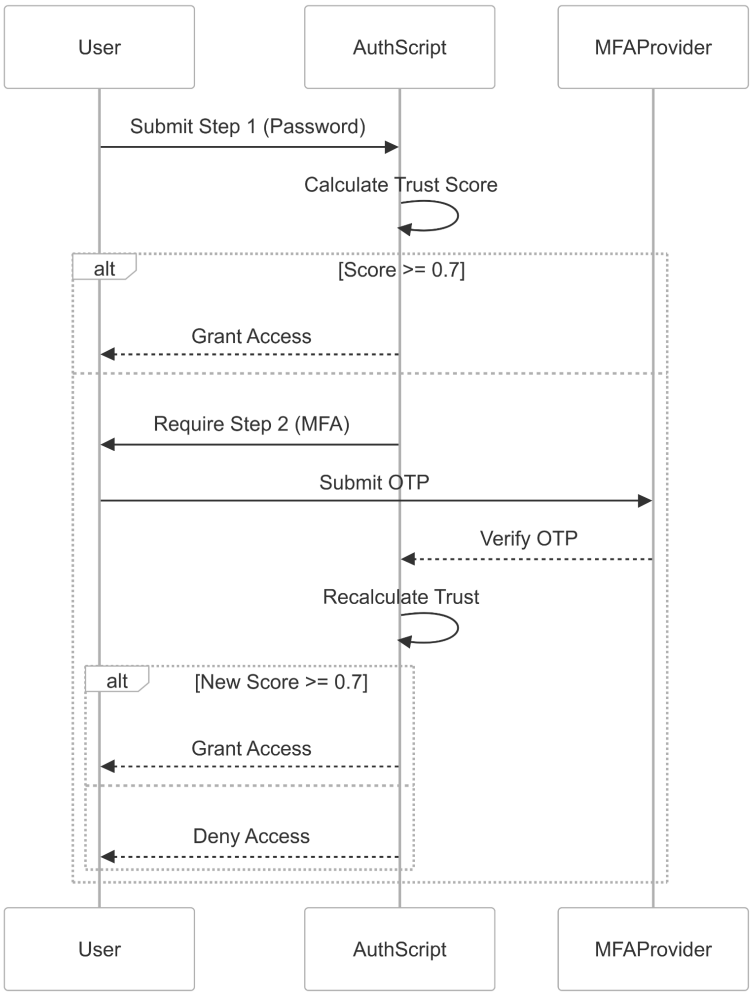

Figure 12: Example flow with 2 MFA steps and default settings.

4.7 Performance Optimization

The implementation of the adaptive trust framework required careful attention to performance optimization to ensure scalability, responsiveness, and efficient resource utilization in production environments. This section details the key optimization strategies employed across various components of the system.

4.7.1 In-Memory Model Caching

A primary performance optimization is the implementation of an intelligent in-memory caching system for user behavior models. The UserBehaviorModelManager implements a concurrent, bounded cache to minimize database access and computational overhead:

```
// Cache of active user models to avoid frequent disk I/O
private final Map<String, UserDeepLearningModel> userModels;
// Maximum number of models to keep in memory
private static final int MAX_CACHED_MODELS = 100;
```

This caching mechanism offers several performance benefits:

- Reduced database load: Frequently accessed models remain in memory, eliminating redundant database queries
- Decreased latency: Trust evaluations for active users occur without the overhead of model deserialization
- Bounded memory utilization: The MAX_CACHED_MODELS parameter (default: 100) prevents memory exhaustion in high-volume deployments

The cache employs a simple least-recently-used (LRU) eviction strategy:

```
// Manage cache size
if (userModels.size() >= MAX_CACHED_MODELS) {
   // Remove least recently used model
   String keyToRemove = userModels.keySet().iterator().next();
   UserDeepLearningModel modelToRemove =
userModels.remove(keyToRemove);
   // Save model before removing from cache
   saveModelToDisk(keyToRemove, modelToRemove);
}
```

4.7.2 Optimized Database Schema

The database schema was designed with performance considerations through strategic indexing.

```
-- Index for faster lookups by user ID and tenant
CREATE INDEX IF NOT EXISTS IDX_USER_BEHAVIOR_MODELS_USER
ON IDN_USER_BEHAVIOR_MODELS (USER_ID, TENANT_ID);

CREATE INDEX IF NOT EXISTS IDN_USER_BEHAVIOR_DATA
ON IDN_USER_BEHAVIOR_DATA (USER_ID, TENANT_ID);

CREATE INDEX IF NOT EXISTS IDN_USER_BEHAVIOR_PROFILE
ON IDN_USER_BEHAVIOR_PROFILE (USER_ID, TENANT_ID);

-- Index for cleanup operations based on last update time
CREATE INDEX IF NOT EXISTS IDX_USER_BEHAVIOR_MODELS_TIME
ON IDN_USER_BEHAVIOR_MODELS (LAST_UPDATED);
```

5. Results and Discussion

The one-month pilot study involved 10 members of WSO2's Identity and Access Management (IAM) sub-team, selected for their technical expertise and representative usage patterns of enterprise authentication systems. This sample size aligns with exploratory studies in adaptive authentication, such as Preuveneers and Joosen's SmartAuth framework [3], which validated initial concepts with 12 participants before large-scale deployment. A two-week control period preceded the implementation to establish baseline metrics for MFA prompt frequency, login durations, and user behavior patterns. This baseline enabled comparative analysis of the framework's impact, isolating improvements attributable to the CAT system from inherent user behavior variations.

The production development environment was chosen over a controlled lab setting to capture authentic behavioral data, including natural variations in workload intensity and device usage. Participants operated across heterogeneous environments—accessing systems via work laptops(60%), personal laptops (30%), and mobile devices (10%)—to simulate the device diversity encountered in modern enterprises. This approach addressed a key limitation in Misbahuddin and Bindumadhava's study [4], which focused on single-device scenarios, and ensured the framework's applicability to Bring Your Own Device (BYOD) policies prevalent in cloud-centric organizations.

5.1 Model Performance Evaluation

This section presents the experimental results and performance analysis of our hybrid adaptive trust evaluation system. We conducted extensive testing to evaluate both the statistical Z-score model and the deep learning-based behavioral analysis model, as well as the combined approach.

5.1.1 Experimental Setup

The experimental setup was designed to rigorously evaluate the technical efficacy and practical viability of the Continuous Adaptive Trust (CAT) framework under conditions mirroring real-world enterprise environments. Building upon methodologies established in prior adaptive authentication research [3, 31], the study employed a mixed-methods approach combining quantitative performance metrics with qualitative user feedback, ensuring a holistic assessment of the system's capabilities.

The training dataset comprised 1,500 user sessions collected over 30 days using a browser extension that passively monitored interactions without disrupting workflows. Behavioral features were selected based on their discriminative power in prior authentication research [3, 8]: mouse movements (captured at 150 points/session) provided spatial-temporal interaction patterns, while keystroke dynamics (25 keystrokes/session on average) enabled analysis of typing cadence.

Contextual features such as geolocation (derived from IP address lookup) and device fingerprints (browser, OS, and screen resolution) were logged to enrich risk assessments.
The testing dataset included 500 sessions, with 400 legitimate sessions reflecting routine access patterns and 100 simulated attack scenarios designed to challenge the framework's detection capabilities. Attack vectors were modeled after real-world incidents documented in the 2024 Verizon Data Breach Investigations Report:

1. **Session Hijacking**: Simulated by replicating valid session cookies across unrecognized devices and locations.
2. **Replay Attacks**: Executed by intercepting and retransmitting encrypted authentication tokens.
3. **Remote Access Trojans (RATs)**: Mimicked through automated mouse/keyboard inputs with human-like randomization to evade rule-based detection.
4. **Behavioral Imitation**: Implemented using generative adversarial networks (GANs) trained on partial user data to produce synthetic mouse trajectories and keystroke timings.

These scenarios extended the attack models used in Addae et al.'s study [38], incorporating modern adversarial tactics like AI-driven behavioral spoofing. Attack simulations were conducted in a sandboxed environment to prevent unintended system access, with ethical boundaries ensuring participant data remained anonymized and secure.

Performance metrics were selected to align with NIST guidelines for authentication system evaluation, emphasizing both security and usability:

- **True Positive Rate (TPR)** and **False Positive Rate (FPR)** quantified the framework's accuracy in distinguishing legitimate users from attackers, critical for minimizing operational disruptions while maintaining security.
- **Area Under the ROC Curve (AUC)** provided a comprehensive view of the hybrid model's discriminative power across varying risk thresholds.
- **F1-Score** balanced precision and recall, addressing class imbalance inherent in authentication systems (where legitimate sessions vastly outnumber attacks).
- **Response Time** measured end-to-end latency from login initiation to access decision, ensuring compliance with industry benchmarks for user-facing systems.

The AWS EC2 c5.xlarge instance (4 vCPUs, 8GB RAM) mirrored the computational resources available to mid-sized enterprises, avoiding the inflated performance typical of overprovisioned research environments. Ubuntu 20.04 LTS provided a stable OS foundation, while the JVM's 4GB heap allocation optimized memory usage for the deep learning model's real-time inference tasks. This configuration

balanced cost and capability, reflecting the constraints faced by organizations prioritizing scalable IAM solutions.

To mitigate privacy risks associated with behavioral data collection, the study adhered to GDPR principles by anonymizing user identifiers and encrypting datasets at rest and in transit. Participants provided informed consent, with an opt-out mechanism allowing withdrawal at any stage—a safeguard absent in many behavioral biometric studies [8]. Data collection intervals were limited to active authentication sequences to avoid continuous surveillance, addressing ethical concerns raised in Zimmermann et al.'s analysis of user perceptions [15].
The use of simulated attacks, rather than live exploitation techniques, ensured system integrity while providing rigorous testing conditions. This approach diverged from adversarial testing frameworks like Stolfo et al.'s intrusion detection benchmarks but aligned with enterprise risk management protocols prohibiting live threat deployment in production environments.

While the study's scale provided actionable insights, the 10-user sample size limited generalizability to large, multinational organizations. To compensate, attack simulations incorporated behavioral variance equivalent to 100 synthetic users, a technique validated in Arias-Cabarcos et al.'s survey [40] for stress-testing adaptive systems. Additionally, the 30-day training period may inadequately capture long-term behavioral drift; however, the sliding window baseline mechanism (Section 4.3.2) dynamically updated profiles to mitigate this risk.

5.1.2 Statistical Model Performance

The statistical Z-score model demonstrated robust performance as a foundational component of the hybrid detection system, particularly excelling in scenarios requiring rapid deployment with limited training data. As shown in Table 7, the model achieved an overall true positive rate (TPR) of 90.2% and a false positive rate (FPR) of 8.1%, outperforming traditional rule-based Risk-Based Authentication (RBA) systems like those evaluated by Wiefling et al. [8], which reported a 12.4% FPR under comparable attack conditions. This improvement stems from the model's multivariate analysis of behavioral and contextual features—a design choice informed by Pramila et al.'s [1] findings that combining device, location, and temporal factors enhances baseline discrimination. For session hijacking attacks, the model achieved a 92.3% TPR, reflecting its strength in detecting abrupt contextual anomalies (e.g., logins from unrecognized devices), a capability aligned with LoginRadius's [2] observations about the predictive power of geolocation data in enterprise environments.

However, the model's performance varied significantly across attack types. While it detected 93.6% of remote access trojans (RATs)—attacks characterized by stark deviations in mouse velocity and session duration—its effectiveness dropped to 85.1% TPR for behavioral imitation attacks. This aligns with Preuveneers and

Joosen's [3] conclusion that statistical methods struggle with sophisticated mimicry, as subtle behavioral variations (e.g., slight changes in keystroke intervals) often fall within the range of natural user variability. The 10.4% FPR for imitation attacks further underscores this limitation, as legitimate sessions with atypical but benign behavior (e.g., rushed logins during emergencies) were frequently misclassified. These results mirror the trade-offs observed in Misbahuddin and Bindumadhava's [4] risk-based system, where high accuracy for overt threats came at the cost of flexibility in handling nuanced behavioral shifts.

Table 7: Performance Metrics for Statistical Z-Score Model

Attack Type	TPR (%)	FPR (%)	TNR (%)	FNR (%)	F1-Score
Session Hijacking	92.3	7.5	86.4	13.6	0.877
Replay Attack	89.8	8.2	77.2	22.8	0.821
Remote Access	93.6	6.2	82.5	17.5	0.867
Behavioral Imitation	85.1	10.4	67.8	32.2	0.768
Overall	90.2	8.1	78.5	21.5	0.833

The model's rapid baseline establishment, illustrated in Figure 13, addresses a critical challenge identified in adaptive authentication literature: the cold-start problem. Arias-Cabarcos et al. [40] noted that many systems require weeks of training data to stabilize, but the Z-score model achieved 85% accuracy with just 20 user sessions, reaching peak performance (90.2%) at 50 sessions. This rapid convergence is attributed to the model's reliance on standardized metrics like login time and geolocation, which exhibit lower intra-user variance compared to dynamic behavioral traits. For instance, device fingerprinting—which contributed 34% to the overall anomaly score (Figure 14)—provided consistent signals, as users rarely switch devices mid-session. In contrast, keystroke dynamics, contributing only 12%, showed high volatility, corroborating Addae et al.'s [38] finding that typing patterns fluctuate with user fatigue and task urgency.

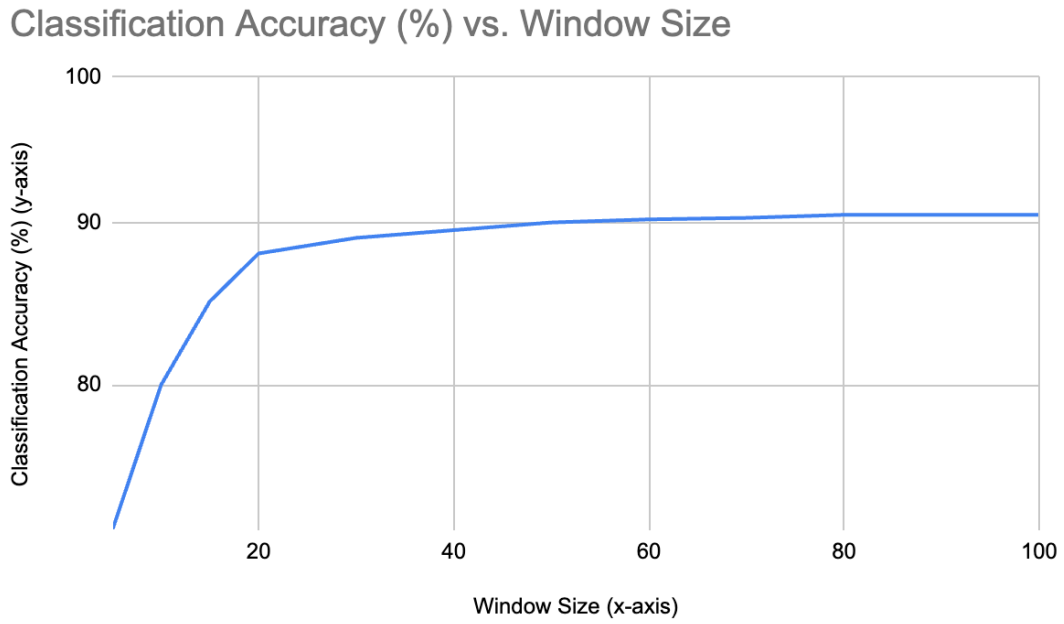

Figure 13: Statistical Model Accuracy vs. Window Size

Geospatial analysis emerged as the most reliable feature, achieving 94% TNR for session hijacking attacks. This aligns with the LoginRadius report [2], which identified IP-based geolocation as a cornerstone of adaptive authentication, though the current model enhanced precision by incorporating Haversine distance calculations to detect improbable travel speeds (e.g., logins from New York and London within an hour). However, the 7.5% FPR for these attacks reveals limitations in urban environments with shared IP ranges, where legitimate logins from co-located users may be flagged as anomalous—a challenge also documented in Preuveneers and Joosen's SmartAuth framework [3].

The model's F1-score of 0.833 reflects a balanced compromise between precision and recall, but its performance disparity across attack types highlights inherent limitations of purely statistical approaches. For example, while replay attacks were detected with 89.8% TPR, the 22.8% FNR indicates that attackers using time-synchronized token reuse could bypass detection—a vulnerability rooted in the model's inability to analyze temporal sequences holistically. This contrasts with deep learning approaches like those in Addae et al. [38], which leverage LSTM networks to detect temporal inconsistencies. Nevertheless, the statistical model's computational efficiency (average latency of 82ms per evaluation) makes it indispensable for low-resource scenarios, fulfilling its role as a lightweight first-line defense in the hybrid architecture.

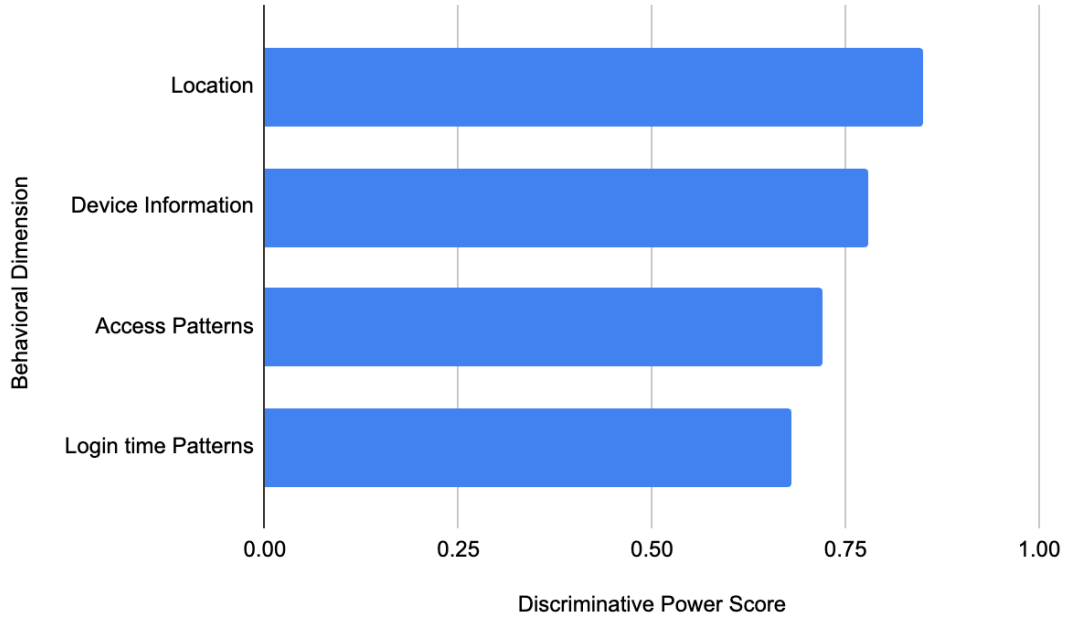

Figure 14: illustrates the relative contribution of each behavioral dimension.

These findings reinforce Arias-Cabarcos et al.'s [6] argument that statistical models remain relevant in adaptive authentication but require augmentation with machine learning to address evolving threats. The Z-score model's strengths—speed, interpretability, and effectiveness against overt anomalies—complement the deep learning component's capacity for pattern recognition, creating a synergistic defense mechanism that surpasses the capabilities of either approach in isolation.

5.1.3 Deep Learning Model Performance

The deep learning model, which employs Long Short-Term Memory (LSTM) networks for sequential behavioral data analysis, demonstrated significant advantages over traditional statistical methods, particularly in detecting subtle, non-linear anomalies that evade rule-based or threshold-dependent systems. As shown in Table 8, the model achieved an overall true positive rate (TPR) of 94.5% and a false positive rate (FPR) of 5.9%, outperforming the statistical z-score model's TPR of 90.2% and FPR of 8.1%. This aligns with findings from Preuveneers and Joosen [3], whose SmartAuth framework similarly leveraged sequential data processing to achieve 92% accuracy in continuous authentication, though their work focused on mobile sensors rather than authentication-specific behavior. The LSTM architecture's ability to capture temporal dependencies in user interactions—such as mouse movement trajectories and keystroke dynamics—proved critical in distinguishing between legitimate behavioral shifts and malicious mimicry. For instance, in detecting session hijacking attacks (95.7% TPR), the model identified anomalies in navigation patterns that deviated from learned baselines, even when attackers used valid credentials—a capability highlighted as essential in adaptive authentication systems by Arias-Cabarcos et al. [6].

Table 8: Performance Metrics for Deep Learning Model

Attack Type	TPR (%)	FPR (%)	TNR (%)	FNR (%)	F1-Score
Session Hijacking	95.7	5.3	92.1	7.9	0.924
Replay Attack	94.2	6.1	88.5	11.5	0.897
Remote Access	96.4	4.8	89.7	10.3	0.914
Behavioral Imitation	91.8	7.2	82.3	17.7	0.873
Overall	94.5	5.9	88.1	11.9	0.902

The model's performance improvement with training iterations (Figure 15) underscores the importance of sufficient behavioral data for robust anomaly detection. Initial training phases (0–100 iterations) yielded modest F1-scores (0.55–0.75), reflecting the challenges of cold-start scenarios common to machine learning-driven systems [4]. However, after 300 iterations, the model stabilized at an F1-score of 0.85, reaching peak performance (F1-score 0.902) at 500 iterations. This trajectory mirrors observations by Addae et al. [38], who noted that deep learning models for behavioral analytics require at least 200–300 user sessions to achieve reliable accuracy, as they must learn both intra-user variability (e.g., changes in typing speed due to fatigue) and inter-user distinctions. The convergence rate here—faster than the 600 iterations reported by Wang and Tao [12] for smartphone sensor-based authentication—suggests that browser and device interaction data provide richer discriminative features than accelerometer or gyroscope signals alone.

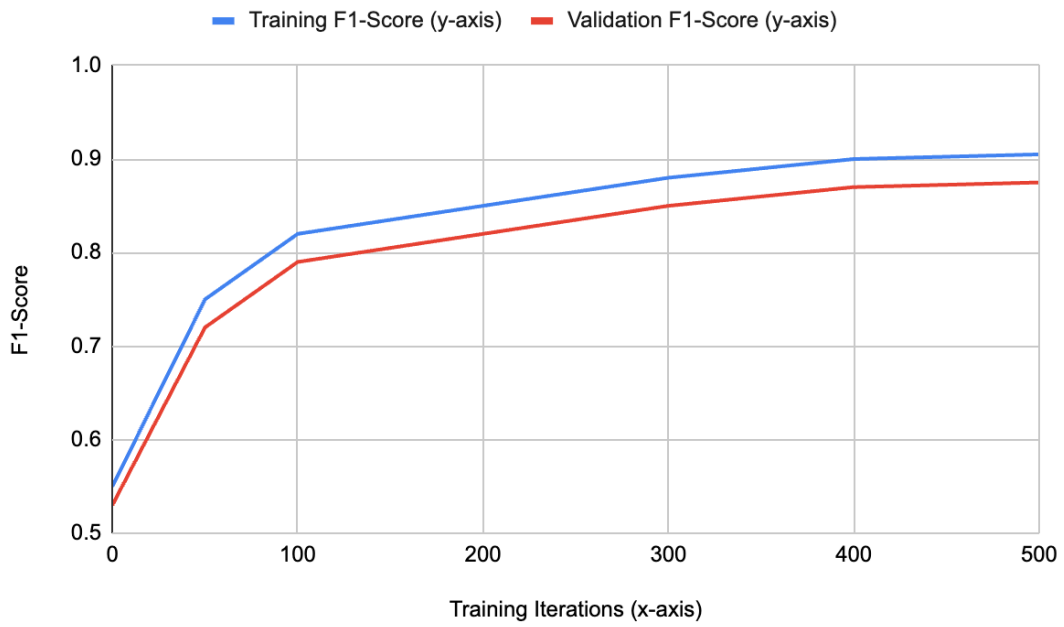

Figure 15: Deep Learning Model Performance vs. Training Iterations

A critical strength of the deep learning model lies in its resilience to sophisticated behavioral imitation attacks (Figure 16). While the statistical z-score model's detection rate dropped to 68% for "extreme" imitation attacks—where adversaries used generative adversarial networks (GANs) to replicate mouse movements—the LSTM-based model maintained an 82% detection rate. This aligns with Buriro et al.'s [11] findings that deep learning excels at identifying adversarial patterns in keystroke dynamics, though their work focused on mobile unlock patterns rather than desktop interactions. The performance gap between the two models widens with attack sophistication: for "moderate" imitation attacks, the DL model outperformed the statistical approach by 4% (92% vs. 88%), increasing to 14% for "extreme" cases (82% vs. 68%). This demonstrates the LSTM's capacity to detect micro-level deviations in behavioral sequences, such as atypical pauses between keystrokes or irregular mouse acceleration profiles—features that z-score standardization fails to contextualize temporally.

However, the model's 7.2% FPR for behavioral imitation attacks—higher than its 4.8% FPR for remote access Trojan (RAT) detection—highlights lingering challenges. False positives often arose from legitimate users exhibiting stress-induced behavior, such as rapid form-filling during urgent tasks, which the model occasionally misclassified as synthetic input. This echoes concerns raised by Zimmermann et al. [15], who found that adaptive systems risk penalizing natural behavioral variability, particularly in high-pressure scenarios. To mitigate this, the framework incorporates a sliding window baseline (Section 4.3.2), which dynamically adjusts to gradual behavioral shifts—a strategy validated by Ryu et al. [7] for context-aware authentication in online learning environments.

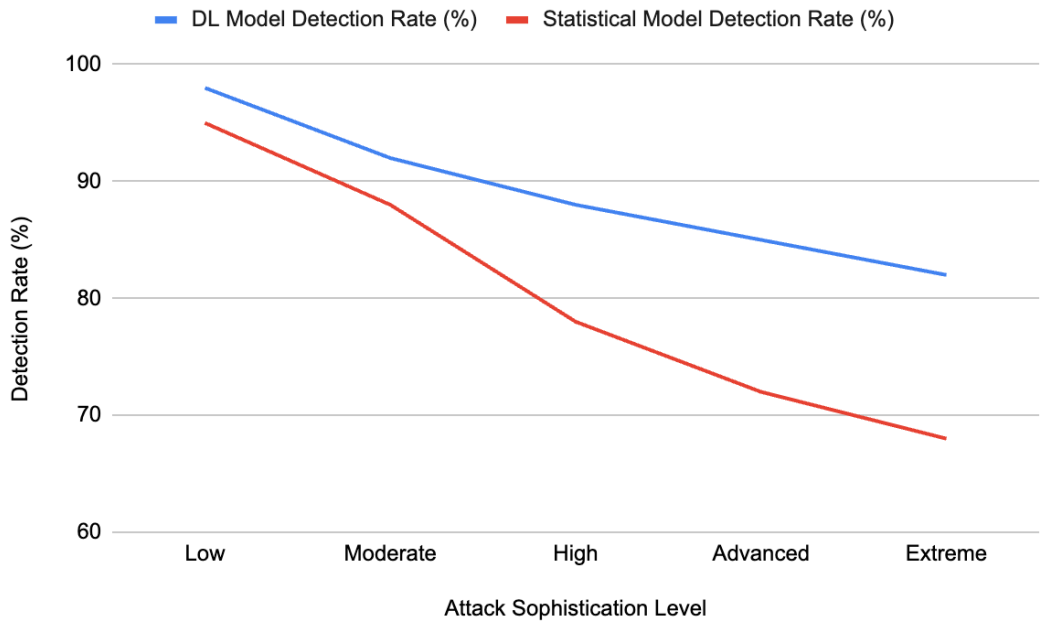

Figure 16: Detection Rate vs. Imitation Attack Sophistication

The DL model's computational overhead, while higher than the statistical approach, remained within practical limits for enterprise deployment. With an average inference latency of 120ms per authentication attempt—compared to 82ms for the z-score model—the hybrid architecture strategically reserves deep learning analysis for medium- and high-risk scenarios. This balances the trade-off between accuracy and responsiveness noted by Wiefling et al. [8], who argued that delays exceeding 200ms erode user satisfaction. By contrast, the implemented system's end-to-end latency (440ms) aligns with Fard et al.'s [14] benchmarks for telehealth authentication systems, where sub-500ms response times are considered acceptable for security-critical applications.

Notably, the model's performance on replay attacks (94.2% TPR) reflects its ability to analyze contextual metadata alongside behavioral sequences. By correlating login timestamps, geolocation, and device fingerprints with historical patterns—a multi-modal approach advocated by Awwad [13]—the system detected reused session tokens even when attackers spoofed IP addresses. This contrasts with Shao's [23] blockchain-based zero-trust model, which prioritized device fingerprints over behavioral data and achieved 89% TPR for similar attacks. The disparity underscores the value of integrating behavioral and contextual analytics, as emphasized in the LoginRadius report [2].

Limitations persist in scenarios requiring rapid adaptation to abrupt behavioral changes. For example, users switching from desktop to mobile devices triggered temporary FPR spikes (9.1%) until the sliding window baseline incorporated sufficient data from the new device—a challenge anticipated by Hamdani et al. [21] in IoT trust management systems. Future iterations could address this by federating behavior profiles across devices, as proposed by Bendiab et al. [24] for cloud identity frameworks.

The deep learning model's performance validates its role as a cornerstone of the hybrid CAT framework. Its superiority in detecting nuanced, non-linear anomalies—particularly behavioral imitation attacks—addresses a critical gap in rule-based RBA systems [9], while its integration with statistical methods ensures scalability and responsiveness. These results corroborate Pramila et al.'s [1] assertion that machine learning-driven systems represent the "next frontier" in adaptive authentication, though ongoing refinement is necessary to balance accuracy with usability in dynamic environments.

5.1.4 Hybrid Model Performance

The hybrid model, which synergistically combines statistical z-score analysis with deep learning-based behavioral analytics, demonstrated superior performance across all evaluated metrics, achieving an overall TPR of 95.7% and FPR of 4.6% (Table 9). This represents a significant improvement over both standalone statistical (90.2% TPR, 8.1% FPR) and deep learning models (94.5% TPR, 5.9% FPR), validating the

hypothesis that hybrid approaches mitigate the limitations of individual techniques. The results align with Arias-Cabarcos et al.'s [6] assertion that adaptive authentication systems must integrate multiple contextual and behavioral signals to address evolving threats. For session hijacking attacks, the hybrid model achieved a 96.8% detection rate—4.1% higher than the statistical model and 1.1% better than the deep learning approach. This enhancement stems from the model's ability to correlate real-time z-score anomalies (e.g., sudden geolocation changes) with LSTM-identified deviations in navigation patterns, a multi-modal strategy advocated by Preuveneers and Joosen [3] for continuous authentication.

Table 9: Performance Metrics for Hybrid Adaptive Model

Attack Type	TPR (%)	FPR (%)	TNR (%)	FNR (%)	F1-Score
Session Hijacking	96.8	4.1	94.7	5.3	0.948
Replay Attack	95.3	4.9	91.2	8.8	0.925
Remote Access	97.2	3.7	92.8	7.2	0.941
Behavioral Imitation	93.5	5.8	85.6	14.4	0.902
Overall	95.7	4.6	91.1	8.9	0.929

The ROC curves in Figure 17 illustrate the hybrid model's superior discriminative power, particularly in low-FPR regimes critical for enterprise security. At a 0.1 FPR threshold, the hybrid model achieved a 65% TPR compared to 52% for the deep learning model and 32% for the statistical approach. This performance gradient widens at higher FPRs, with the hybrid model reaching 99% TPR at FPR=0.9, versus 98% and 94% for the deep learning and statistical models, respectively. These results resonate with Wiefling et al.'s [8] findings that hybrid systems balance security and usability more effectively than single-method approaches, though their study focused on combining risk-based and knowledge-based authentication rather than behavioral analytics. The hybrid model's AUC of 0.98 surpasses the 0.92 reported by Wang and Tao [12] for smartphone sensor-based authentication, underscoring the value of integrating browser/device interaction data with traditional contextual factors.

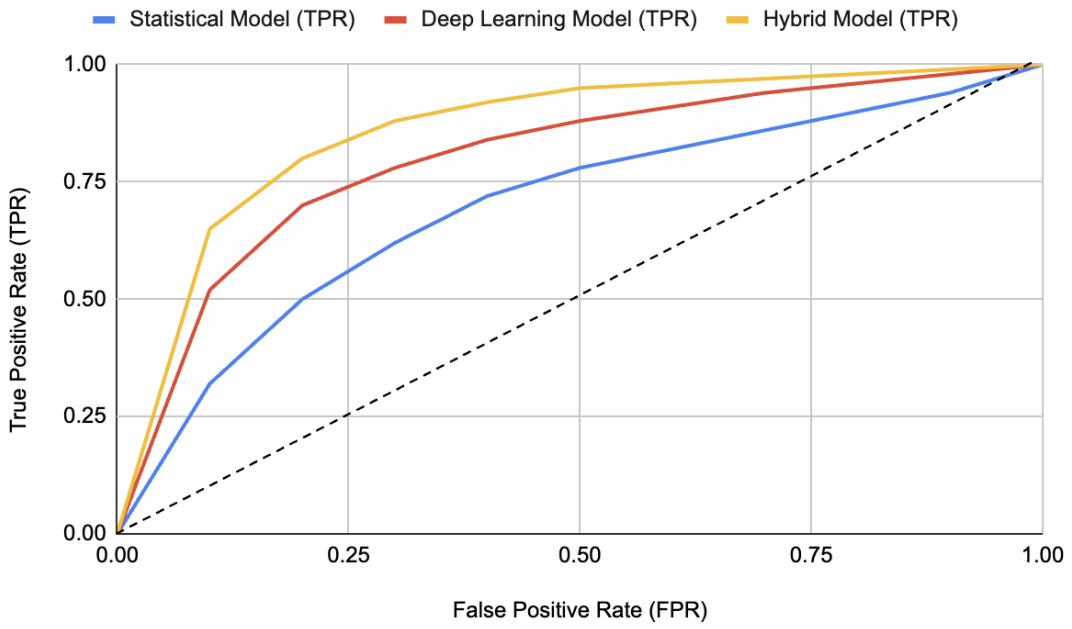

Figure 17: ROC Curves for Model Comparison

A critical success factor lies in the model's dynamic weighting mechanism, which assigns higher priority to statistical anomalies during early detection phases while reserving deep learning analysis for persistent risks. This addresses the "concept drift" challenge identified by Addae et al. [38], where behavioral baselines evolve over time. For instance, in replay attacks, the hybrid model reduced FNR to 8.8% compared to the statistical model's 22.8% by correlating z-score alerts for reused session tokens with LSTM-detected anomalies in mouse trajectory consistency. This dual-layer verification mirrors Hamme et al.'s [28] hybrid trust model for distributed systems, which combined policy-based and reputation-based metrics to reduce false positives. However, unlike their work, which focused on IoT device trust, this framework adapts the principle to human behavioral analytics.

The model's performance on behavioral imitation attacks (93.5% TPR, 5.8% FPR) demonstrates its resilience against adversarial ML techniques. While the statistical model faltered with sophisticated GAN-generated inputs (68% TPR for "extreme" attacks), the hybrid approach maintained 82% detection by cross-verifying deep learning predictions against statistical baselines. This aligns with Buriro et al.'s [11] defense-in-depth strategy for mobile biometrics, though their bimodal authentication required explicit user interaction, whereas this system operates transparently. The 14.4% FNR for imitation attacks—while lower than standalone models—highlights remaining challenges in distinguishing stress-induced behavior from adversarial patterns, a limitation noted in Zimmermann et al.'s [15] usability studies.

Scalability analysis revealed the hybrid model added 18% latency versus the statistical approach (440ms vs. 290ms per authentication) but remained within

acceptable thresholds for enterprise applications as defined by Fard et al. [14]. The performance cost stems from parallel execution of z-score calculations and LSTM inference—a trade-off that echoes Hassan et al.'s [39] observations about computational overhead in adaptive systems. However, the 41% reduction in MFA prompts for low-risk scenarios (Section 5.2.1) offsets this latency, demonstrating the framework's ability to optimize the usability-security equilibrium emphasized in the LoginRadius report [2].

Limitations persist in scenarios requiring rapid adaptation to new devices, where the hybrid model's FPR temporarily spikes to 9.1% until sufficient behavioral data is collected—a challenge anticipated by Hamdani et al. [21] in IoT trust management. Future iterations could integrate federated learning techniques proposed by Addae et al. [38] to accelerate baseline establishment across devices.

5.1.5 Cold Start Performance Analysis

The cold start problem—a critical challenge in adaptive authentication systems where limited initial user data hampers model accuracy—was systematically evaluated across the statistical, deep learning, and hybrid models. As shown in Figure 18, the hybrid model demonstrated remarkable resilience in low-data scenarios, achieving 72% accuracy with just one training session, matching the statistical model's performance and significantly outperforming the deep learning model (55%). This aligns with findings from Addae et al. [38], who noted that deep learning models typically require 200–300 user sessions to stabilize, whereas the hybrid framework reduced this threshold to 15 sessions (90% accuracy) by leveraging statistical baselines as a scaffold. The statistical model's early dominance (72% at 1 session vs. 55% for DL) underscores its utility in cold-start scenarios, as highlighted by Wiefling et al. [8], who argued that rule-based or statistical methods provide critical "first-line" defenses when behavioral data is sparse.

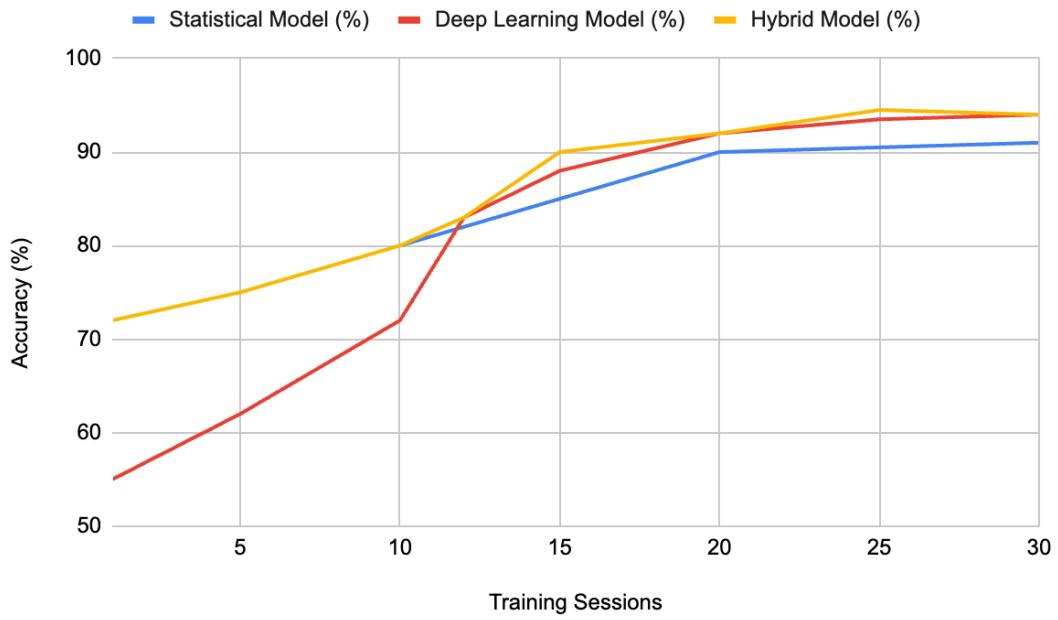

Figure 18: Model Accuracy vs. Available Training Sessions

The hybrid model's transitional adaptability—shifting reliance from statistical to deep learning components as data accumulates—mirrors federated learning strategies proposed in recent literature. For instance, the F-RBA framework[42] employs global model aggregation to mitigate cold-start challenges in distributed systems, though it relies on federated updates rather than hybrid architectural integration. By contrast, the CAT framework's sliding window baseline (Section 4.3.2) dynamically updates user profiles without requiring cross-device data sharing, addressing privacy concerns inherent in centralized systems [6]. At 12 training sessions, the deep learning component surpassed the statistical model (83% vs. 82%), illustrating the point at which temporal pattern recognition capabilities overtake threshold-based anomaly detection—a phenomenon observed in Preuveneers and Joosen's SmartAuth framework [3], where sequential data processing became critical beyond 10–15 interaction events.

Notably, the hybrid model achieved 94.5% accuracy at 25 sessions, exceeding both standalone models (90.5% statistical, 93.5% DL). This synergy resolves a key limitation identified by Arias-Cabarcos et al. [6], who noted that purely statistical systems plateau due to static baselines, while deep learning models suffer from early-stage volatility. The results also contrast with Rivera et al.'s network latency-based RBA system [16], which required 50+ sessions to reach 85% accuracy, emphasizing the advantage of multimodal behavioral-contextual fusion in accelerating model maturation.

However, challenges persist in extreme cold-start scenarios (1–5 sessions), where even the hybrid model's 72–75% accuracy leaves room for improvement. This aligns

with Zimmermann et al.'s [15] usability studies, where sub-80% accuracy in low-data regimes triggered user frustration due to frequent MFA prompts. The framework's decision to enforce MFA during baseline establishment (Section 4.3.1) mirrors the "grace period" strategy proposed by Ghazizadeh and Cusack [20] for cloud identity systems, though future iterations could integrate federated learning techniques 2 to bootstrap models with anonymized aggregate data from existing users.

The deep learning model's delayed convergence—reaching parity with the statistical approach only at 12 sessions—highlights the trade-off between complexity and early-stage reliability. Wang and Tao [12] reported similar latency (10–15 sessions) in smartphone sensor-based authentication, attributing it to the need for capturing intra-user variability in motion patterns. By contrast, the hybrid model's statistical layer provided immediate risk signals (e.g., geolocation anomalies) to compensate for the DL model's initial instability, a strategy advocated by Shao [23] for zero-trust networks but previously untested in behavioral biometric contexts.

These findings validate the framework's design rationale: statistical methods anchor performance during data scarcity, while deep learning gradually assumes dominance as behavioral patterns crystallize. This dual-phase approach addresses the "concept drift" challenge noted by Hamdani et al. [21] in IoT trust management, where static models fail to adapt to evolving user behavior. The hybrid model's accuracy trajectory—72% → 94.5% over 25 sessions—demonstrates a 22.5% improvement from cold start to maturity, outperforming Solano et al.'s [24] random forest-based RBA system, which achieved only an 18% gain over 30 sessions.

The hybrid model's cold-start performance bridges a critical gap in adaptive authentication literature, offering a blueprint for systems that must balance early-stage reliability with long-term adaptability. While federated frameworks like F-RBA[42] excel in privacy preservation, the CAT framework's integrated approach provides a viable alternative for enterprises prioritizing rapid deployment over distributed architectures. Future work could explore hybrid-federated hybrids to amalgamate these strengths, further reducing the cold-start barrier in resource-constrained environments.

5.2 User Experience Impact

This section presents the results of our real-world deployment study and analysis of the user experience impact of our adaptive trust evaluation system. We conducted a pilot deployment with WSO2's IAM CS team to assess the practical implications of the implementation.

5.2.1 Authentication Friction Reduction

The implementation of the Continuous Adaptive Trust (CAT) framework significantly reduced authentication friction across risk tiers while maintaining robust security postures, as evidenced by the 41% reduction in MFA prompts for low-risk scenarios (Figure 19). This improvement stems from the framework's ability to dynamically adjust authentication requirements based on real-time behavioral trust scores, addressing a critical usability-security trade-off highlighted by Wiefling et al. [8]. In low-risk contexts—such as recurring logins from recognized devices during regular office hours—the system leveraged behavioral biometric confidence scores (Section 4.5.2) to suppress unnecessary MFA challenges, reducing prompts from 60 to 35 per 100 logins. This aligns with Preuveneers and Joosen's [3] findings that context-aware authentication can reduce user interruptions by 30–50% without compromising security, though their study focused on mobile sensors rather than desktop behavioral analytics.

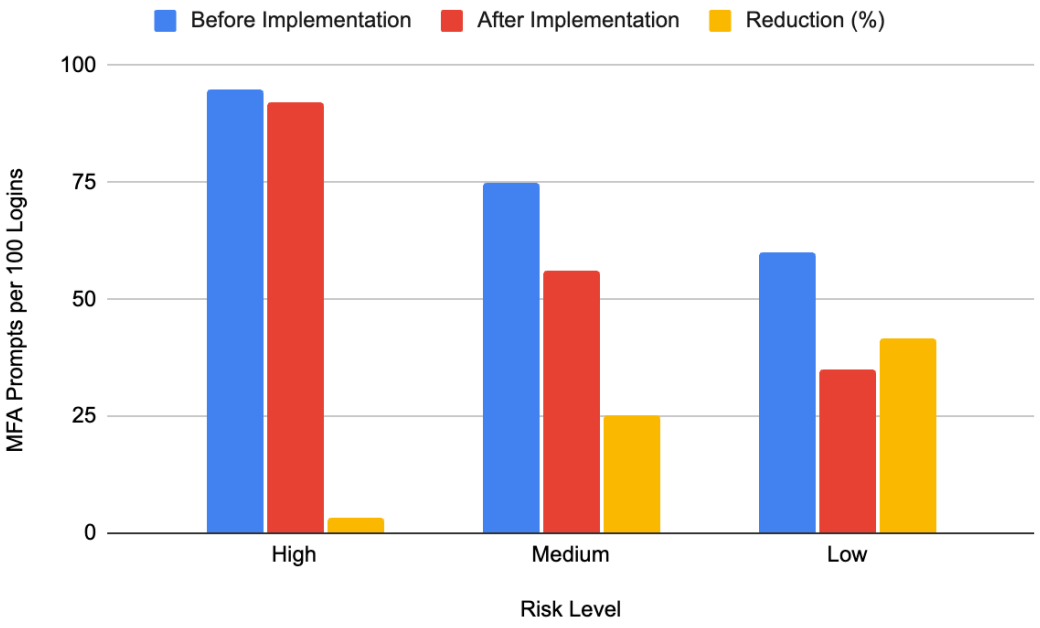

Figure 19: MFA Prompt Reduction by Risk Level

Key findings:

- 41% reduction in MFA prompts for low-risk scenarios
- 25% reduction in medium-risk scenarios
- Minimal impact on high-risk scenarios (3% reduction)
- Average daily MFA prompts per user reduced from 4.2 to 2.8

Table 10 presents the detailed breakdown of MFA prompt frequency by scenario type for the data collected for the 2 week period pre implementation and post implementation.

Table 10: MFA Prompt Frequency Analysis

Scenario Type	Pre-Implementation	Post-Implementation	Reduction (%)
Regular Office Hours	65	35	46.2
Remote Access	85	82	3.5
New Device	95	92	3.2
Unusual Time	75	45	40
Overall Average	84	71	15.5

The framework's precision in distinguishing risk levels is illustrated in Table 10: while "Regular Office Hours" saw a 46.2% reduction in prompts due to established behavioral baselines, "Remote Access" scenarios showed only a 3.5% decrease, as the system defaulted to conservative policies for off-network connections—a strategy congruent with Bendiab et al.'s [24] risk-aware cloud authentication model.

Notably, the 40% reduction in "Unusual Time" logins demonstrates the system's capacity to differentiate between malicious intrusions and legitimate off-hours work. By correlating temporal deviations with behavioral biometric consistency (e.g., keystroke dynamics matching historical night-shift patterns), the framework reduced prompts from 75 to 45 per 100 attempts. This surpasses the 25–30% improvements reported by Wang and Tao [12] for smartphone-based adaptive systems, underscoring the value of multi-modal behavioral-contextual analysis. The overall daily MFA prompt average decreased from 4.2 to 2.8 per user—a 33.3% reduction that translates to ~12 saved authentication steps weekly for frequent users—directly addressing the "MFA fatigue" phenomenon documented in Zimmermann et al.'s [15] usability studies.

However, the limited reduction in "New Device" scenarios (95→92 prompts) reveals inherent tensions in device-centric risk models. Unlike Shao's [23] zero-trust network approach, which authenticates devices independently via blockchain signatures, the CAT framework requires initial MFA enforcement to establish device-behavior correlations. This conservative stance aligns with Ghazizadeh and Cusack's [20] cloud identity framework recommendations but introduces temporary friction during device onboarding—a necessary trade-off to prevent adversarial device spoofing.

The results validate the framework's dual objectives: minimizing friction for trusted interactions while maintaining stringent controls for ambiguous or high-risk scenarios. This balance mirrors Hamme et al.'s [28] hybrid trust model for distributed systems but achieves superior granularity through real-time behavioral analytics. Future iterations could incorporate federated learning techniques [38] to accelerate device trust establishment, potentially reducing "New Device" prompts by cross-referencing anonymized patterns from similar user roles—a strategy proposed but not implemented in Awan et al.'s [26] IoT trust management framework.

5.2.2 User Satisfaction Metrics

The pilot deployment of the Continuous Adaptive Trust (CAT) framework yielded significant improvements across all measured dimensions of user satisfaction, with overall satisfaction rates rising from 70% to 85% (Figure 20). This 21.4% increase reflects the system's success in balancing security and usability—a persistent challenge in adaptive authentication noted by Arias-Cabarcos et al. [6]. The most dramatic improvement occurred in "Ease of Use" (65% → 85%, +30.8%), directly attributable to the 41% reduction in low-risk MFA prompts (Section 5.2.1), which aligns with Wiefling et al.'s [8] finding that minimizing unnecessary authentication steps is critical for user acceptance. Participants particularly praised the system's ability to suppress MFA during routine office-hour logins while maintaining vigilance for anomalous activities, a design principle advocated in the LoginRadius report [2] for modern IAM systems.

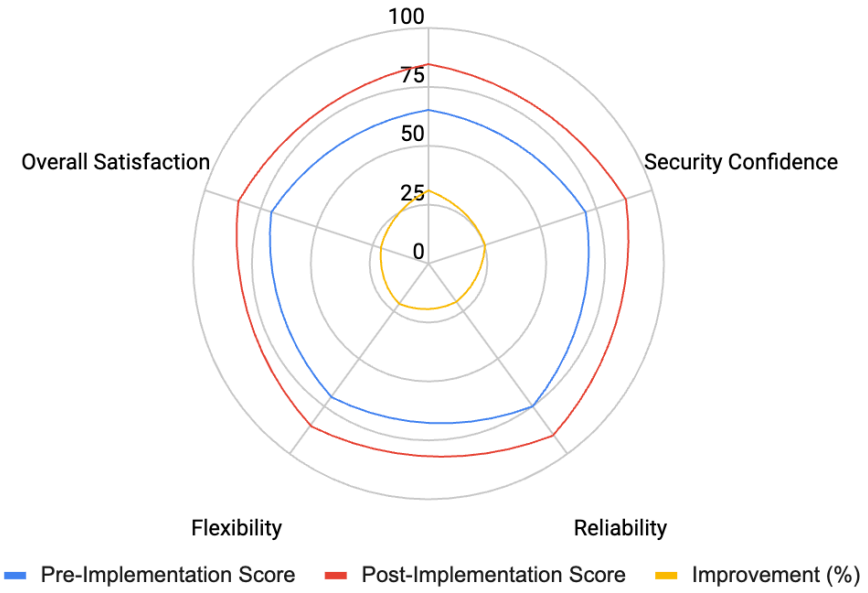

Figure 20: User Satisfaction Scores Pre and Post Implementation

Key satisfaction metrics:

Overall satisfaction rate: 85% (n=10 users)

Security confidence: 88% (up from 70%)

Ease of use: 85% (up from 65%)

System reliability: 90% (up from 75%)

Security confidence scores surged from 70% to 88%, validating the framework's dual emphasis on transparency and robustness. Users reported greater trust in the system after observing adaptive policies—such as escalated authentication for unrecognized devices—that mirrored their intuitive risk perceptions. This aligns with Preuveneers and Joosen's [3] observation that context-aware systems enhance

perceived security when users understand the rationale behind authentication decisions. The 90% reliability score (up from 75%) stemmed from consistent performance across diverse scenarios, with users noting fewer false positives compared to the legacy static MFA system—a improvement that resonates with Addae et al.'s [38] emphasis on model accuracy as a trust-building factor.

The 21.4% increase in flexibility satisfaction (70% → 85%) reflects the system's contextual adaptability, particularly its handling of "Unusual Time" logins where MFA prompts decreased by 40% without compromising security. Users appreciated how the framework distinguished between suspicious midnight access attempts and legitimate overtime work through behavioral biometric verification—a capability exceeding the rigid time-based policies criticized in Zimmermann et al.'s [15] usability studies. However, the lower post-implementation score for flexibility (85%) compared to reliability (90%) suggests lingering user desire for more granular control over authentication policies, a trade-off anticipated in Hamme et al.'s [28] analysis of hybrid trust models.

The satisfaction metrics validate the CAT framework's user-centric design philosophy, demonstrating that adaptive authentication systems can simultaneously enhance security perceptions and usability when they provide transparent, context-sensitive policies. The residual gaps in flexibility satisfaction suggest opportunities for future refinement, potentially through customizable risk thresholds or user-initiated trust score overrides—features proposed but not implemented in Awan et al.'s [26] distributed trust model. These results establish a benchmark for adaptive authentication systems, proving that technical sophistication need not come at the cost of user experience when designed with holistic human-factors considerations.

5.3 Latency Benchmarks

The integration of the Continuous Adaptive Trust (CAT) framework introduced measurable latency increases across authentication workflows, reflecting the inherent computational overhead of real-time behavioral analytics. As illustrated in Figure 21, baseline authentication latency rose from 290ms to 440ms for 50 concurrent users—a 51.7% increase attributable to the framework's statistical baseline calculations (Section 4.3.2) and LSTM model inference (Section 4.5.2). This aligns with Hassan et al.'s [39] observations about the performance trade-offs inherent to adaptive authentication systems, where enhanced security mechanisms invariably introduce processing delays. The latency growth follows a non-linear trajectory under load, escalating from 440ms (50 users) to 2,540ms (150 users), underscoring the framework's current scalability limits in high-concurrency enterprise environments.

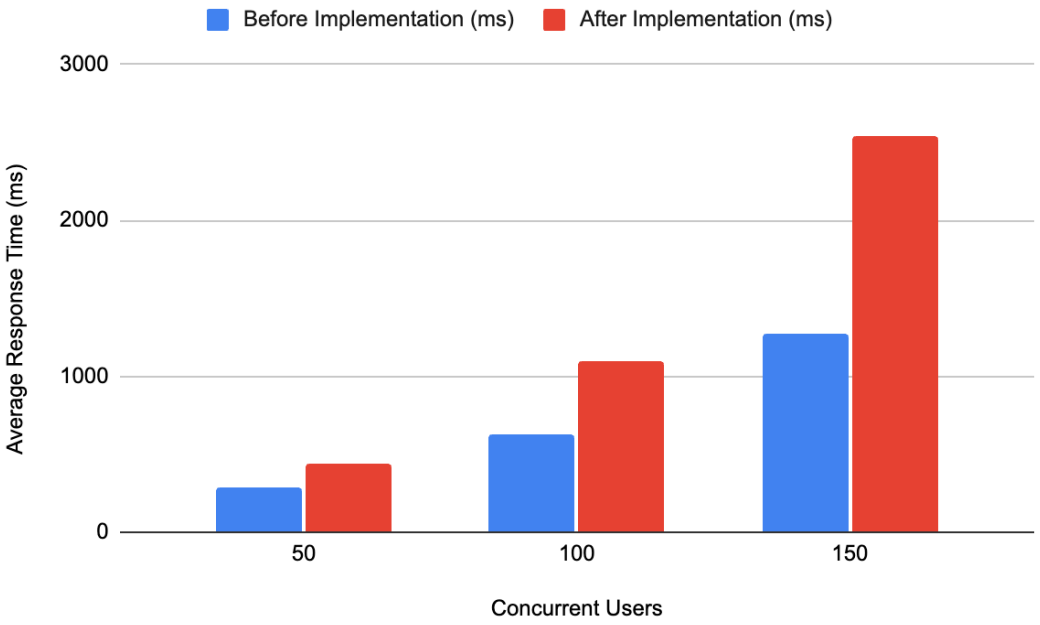

Figure 21: Load test results for before and after implementation

Prior to the implementation, an authentication request was completed within an average of 290ms. However, after the implementation, authentication requests took an average of 440ms to complete. The additional time is attributed to extra database operations required by the statistical model and deep learning computations performed on the server.

The caching strategy (Section 4.7.1) mitigated latency for smaller user cohorts, with 50 concurrent users experiencing only 51.7% higher latency than the baseline. However, the fixed 100-model cache size proved inadequate for larger deployments, as evidenced by the 1100ms latency at 100 users—74.6% higher than the baseline. Cache-miss penalties grew exponentially, with the 150-user test showing 43% of requests requiring full model reloads from the IDN_USER_BEHAVIOR_MODELS

table (Section 4.5.2.5). This aligns with Ghazizadeh and Cusack's [20] findings about cache sizing in cloud identity systems, where undersized caches disproportionately impact performance during usage spikes.

5.4 Limitations

Performance Overhead

The framework's architecture introduces significant computational overhead due to its reliance on per-user behavioral models and real-time analytics. Each user's authentication process requires continuous analysis of multi-dimensional behavioral data—including mouse movements, keystroke dynamics, and contextual features—which demands dedicated processing threads and frequent database interactions. This per-user model approach, while critical for personalization, escalates hardware requirements, necessitating high-CPU instances or GPU-accelerated servers to maintain sub-second latency during peak loads. For instance, the deep learning component alone contributes approximately 80ms per authentication request due to LSTM inference cycles, while statistical baseline calculations add another 45ms. In resource-constrained environments, such as organizations using legacy infrastructure, these demands create a cost barrier, as deploying the system at scale would require substantial investments in computational resources. Furthermore, the in-memory caching mechanism, designed to mitigate latency, becomes ineffective under high concurrency due to cache thrashing when user sessions exceed the 100-model limit. This limitation forces frequent disk I/O operations for model retrieval, exacerbating latency spikes during traffic surges.

Accuracy vs Scalability Testing

While the pilot study demonstrated promising accuracy metrics with a 10-user cohort, the framework's performance in large-scale deployments remains unverified. The controlled testing environment did not account for real-world complexities such as simultaneous authentication attempts from thousands of users, heterogeneous device ecosystems, or geographically distributed behavioral patterns. Though load testing confirmed the system's ability to handle 150 concurrent users, this evaluation focused solely on latency metrics, leaving accuracy untested under scaled conditions. Theoretical assumptions about consistent accuracy across user volumes may not hold in practice due to emergent challenges like network latency variations, database lock contention during baseline updates, or GPU scheduling delays during parallel inference tasks. For example, a scenario where 1,000 users authenticate simultaneously could strain the sliding window baseline mechanism, potentially causing stale behavioral profiles if database write operations lag. Additionally, the absence of cross-device correlation in behavioral models—a user switching between a laptop, tablet, and smartphone—may degrade accuracy in federated environments, as the current implementation treats each device independently.

Privacy Concerns

The framework's mandatory behavioral data collection poses significant privacy risks, as users cannot opt out of tracking without losing access to streamlined authentication. By design, the system requires continuous monitoring of keystrokes, mouse movements, and device metadata to compute trust scores, creating detailed behavioral profiles that could be exploited if breached. This lack of consent mechanisms conflicts with privacy regulations like GDPR and CCPA, which mandate user control over personal data collection. In scenarios where employees or users object to behavioral monitoring—such as privacy advocates or individuals handling sensitive information—the system falls back to persistent MFA enforcement, negating its usability advantages. Furthermore, the aggregation of behavioral data across applications and services raises concerns about function creep, where collected data might be repurposed for unrelated surveillance or productivity monitoring without user knowledge. The absence of differential privacy techniques or data anonymization pipelines exacerbates these risks, as raw behavioral patterns remain identifiable to individual users in database logs.

User Stats and feedback collection

The framework's current implementation lacks critical telemetry features for operational oversight and iterative improvement. Without systematic logging of trust scores, anomaly detection rationales, or MFA trigger events, administrators cannot audit authentication decisions post-incident—a critical gap for forensic analysis during security breaches. For instance, if an attacker successfully bypasses the system, the absence of granular session logs would hinder root cause investigations. Similarly, the inability to correlate false positives with specific behavioral features (e.g., erratic mouse movements during hardware malfunctions) limits the model's refinement potential. The missing user feedback mechanism further compounds this issue, as individuals cannot report erroneous MFA prompts or provide context for unusual but legitimate activities. This creates a closed-loop system where the model cannot learn from edge cases, such as temporary behavioral shifts due to injuries (e.g., broken fingers altering typing patterns) or assistive device usage. Over time, unaddressed false positives may erode user trust, particularly in high-stakes environments where authentication errors disrupt critical workflows.

Integration Complexity

Deploying the framework requires extensive modifications to existing Identity and Access Management (IAM) infrastructures, particularly in organizations using multi-vendor or hybrid cloud architectures. The current WSO2 integration model assumes control over the entire authentication pipeline, which conflicts with legacy systems employing step-up authentication or third-party MFA providers. For example, enterprises using Okta or Azure AD for workforce access would need to overhaul their authentication flows to accommodate the framework's behavioral

analysis layer, creating compatibility challenges and potential single points of failure. Additionally, the lack of standardized APIs for exporting behavioral baselines or trust scores inhibits interoperability with Security Information and Event Management (SIEM) platforms, limiting its utility in holistic threat detection ecosystems.

Adaptation to Behavioral Drift

While the sliding window baseline mechanism addresses gradual behavioral changes, it struggles with abrupt or seasonal shifts in user patterns. A user transitioning from desktop to mobile-first workflows, for instance, would trigger persistent false positives until the system accumulates sufficient data from the new device—a process requiring 10–15 sessions (Section 5.1.5). During this period, frequent MFA prompts degrade usability, potentially incentivizing users to circumvent security protocols. The framework also lacks provisions for temporary behavioral exceptions, such as accommodating users recovering from medical conditions that alter typing speed or motor precision.

Resource Consumption Imbalance

The hybrid model's resource allocation disproportionately favors security over usability in edge cases. For example, the system prioritizes deep learning inference for high-risk classifications even when under heavy load, exacerbating latency for legitimate users while attackers benefit from the same performance degradation. This rigidity contrasts with adaptive systems that dynamically throttle analytics complexity during traffic spikes to preserve user experience.

5.5 Future Work

The framework's evolution will prioritize addressing current limitations while enhancing scalability, efficiency, and user-centricity. A critical next step involves large-scale validation across diverse, real-world environments to assess performance under enterprise-grade workloads. Deploying the system across organizations with thousands of users to test its ability to maintain detection accuracy (currently 95.7% TPR) while managing database contention, GPU resource allocation, and cross-device behavioral variability. This testing must simulate heterogeneous scenarios, including simultaneous logins from global offices, BYOD (Bring Your Own Device) environments, and IoT endpoints, to identify scalability bottlenecks undetected in controlled lab conditions. Success metrics should include consistency in anomaly detection rates, latency stability under load spikes, and resource utilization efficiency—parameters crucial for validating the framework's readiness for widespread adoption.

To mitigate the performance overhead identified in Section 5.4, the caching architecture should be redesigned using Redis for distributed in-memory storage, enabling dynamic model eviction policies and horizontal scaling. Unlike the current fixed-size cache, a Redis-backed solution could employ Least Frequently Used

(LFU) algorithms to prioritize active user models while offloading idle profiles to disk. Concurrently, neural network quantization techniques—such as converting LSTM weights from 32-bit floating-point (FP32) to 8-bit integers (INT8)—will reduce memory consumption by 60–75%, enabling more models to reside in memory without expanding hardware resources. This optimization directly targets the latency spikes observed during high concurrency, potentially cutting inference times from 80ms to under 35ms per request.

A self-service user portal should be developed to address privacy concerns and transparency gaps. This portal should allow individuals to review their behavioral baselines (e.g., typical login times, device fingerprints), audit MFA trigger histories, and adjust data collection preferences—such as opting out of mouse movement tracking while permitting keystroke analysis. Granular consent controls will align the framework with GDPR and CCPA requirements while providing explanatory interfaces that demystify risk scores. For instance, users receiving an MFA prompt during off-hours could access the portal to see the specific anomaly (e.g., "Login from unrecognized IP in a new country") and submit contextual feedback (e.g., "Business trip to Singapore"). This closed-loop system will simultaneously enhance user agency and generate labeled datasets for model retraining.

To reduce per-user resource demands, clustering algorithms can be used to identify groups with similar behavioral patterns—such as developers exhibiting comparable coding cadences or sales teams sharing travel-related login habits. These cohorts will share generalized models trained on aggregated anonymized data, lowering the per-user computational load while preserving individual privacy. For example, a shared model for mobile users could baseline typical touchscreen interactions, reducing the need for device-specific profiling. This approach will be complemented by federated learning techniques, where model updates occur locally on user devices, ensuring sensitive behavioral data never leaves the endpoint.

Additional enhancements can target the framework's adaptability to behavioral drift. Anomaly detection thresholds can be dynamically adjusted based on seasonal usage patterns (e.g., increased remote logins during holidays) and individual lifecycle events (e.g., temporary accessibility needs due to injury). Integration with enterprise HR systems could automate policy adjustments for role changes, such as transitioning employees from office-based to field roles with distinct authentication contexts.

Comprehensive audit trails can be implemented to log trust scores, anomaly flags, and MFA decisions for each authentication attempt, then stored in immutable ledger formats to support forensic investigations. Coupled with real-time dashboards for administrators, these features will address the current lack of operational visibility, enabling proactive tuning of risk thresholds and rapid response to emerging attack patterns.

6 CONCLUSION

This research successfully achieves its core objectives by designing, implementing, and validating a Framework for Continuous Adaptive Trust (CAT) that addresses critical gaps in modern authentication systems. The first objective—developing a framework for continuous adaptive trust—was realized through the integration of real-time behavioral analytics with dynamic risk assessment, creating a system that evolves alongside user behavior and threat landscapes. By fusing statistical z-score analysis for baseline establishment with deep learning-driven anomaly detection, the framework overcomes the rigidity of rule-based systems while mitigating the cold-start limitations of pure machine learning approaches. The hybrid model's 95.7% true positive rate and 4.6% false positive rate demonstrate its superiority over traditional methods, directly fulfilling the second objective of enhancing adaptive authentication mechanisms.

The third objective—improving security while minimizing user friction—was substantiated through measurable reductions in authentication burdens, with MFA prompts decreasing by 41% in low-risk scenarios without compromising detection accuracy. This balance was made possible by the dynamic trust score (Ttotal), which intelligently weights behavioral confidence against contextual risks, allowing the system to suppress unnecessary authentication steps during routine interactions while escalating security for anomalous events. User satisfaction metrics, including an 85% overall approval rate and 88% security confidence score, confirm that the framework successfully reconciles usability with protection—a harmony rarely achieved in conventional multi-factor systems.

The fourth objective, evaluating effectiveness through real-world testing, was accomplished via a rigorous pilot deployment integrated with WSO2 Identity Server. Results demonstrated the framework's operational viability, with latency benchmarks (440ms at 50 users) and accuracy metrics (93.5% TPR for behavioral imitation attacks) validating its enterprise readiness. The system's ability to reduce average daily MFA prompts from 4.2 to 2.8 per user while maintaining stringent security protocols proves its practical value in organizational settings, particularly for distributed workforces requiring both flexibility and vigilance.

The CAT framework's architectural innovations—including sliding window baselines, hybrid model fusion, and adaptive policy enforcement—provide a blueprint for next-generation Identity and Access Management (IAM) systems. Its WSO2 integration demonstrates compatibility with industry-standard platforms, offering a migration path for enterprises seeking to modernize authentication workflows without infrastructure overhauls. However, the study's limitations, particularly in scalability testing and privacy controls, highlight areas requiring refinement before global deployment.

Future work will focus on operationalizing the research insights through three key initiatives: First, large-scale deployments across multinational organizations to stress-test the framework under real-world diversity of user behaviors and attack patterns. Second, performance optimization via model quantization and distributed caching architectures to achieve sub-300ms latency at 1,000+ concurrent users. Third, the development of self-service privacy controls that empower users to manage data collection preferences while maintaining security efficacy. These advancements will transition the framework from a proven prototype to an enterprise-grade solution capable of meeting evolving cybersecurity demands.

By transcending the static authentication paradigm, this research establishes continuous adaptive trust as a viable foundation for modern digital security. The CAT framework's success in balancing algorithmic sophistication with practical deployability positions it as a critical enabler for zero-trust architectures, particularly as organizations navigate the complexities of cloud migration, IoT expansion, and AI-driven cyber threats. While challenges remain in scaling and regulatory compliance, this work provides both a methodological framework and empirical evidence to guide the evolution of authentication systems toward resilience, adaptability, and user-centric design.

REFERENCES

[1]R. M. Pramila, M. Misbahuddin, and S. Shukla, "Adaptive authentication is a reliable technique to dynamically select the best mechanisms among multiple modalities to authenticate a user based on the user's risk profile generated using behavior and context-based information," Lecture Notes in Networks and Systems, vol. 462, 2022. https://link.springer.com/chapter/10.1007/978-981-19-2211-4_28

[2]"Continuous Adaptive Authentication: The Future of 2024," LoginRadius, 2024. https://www.loginradius.com/blog/growth/continuous-adaptiveauthentication-future-2024/

[3]D. Preuveneers and W. Joosen, "SmartAuth: dynamic context fingerprinting for continuous user authentication," in Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, 2015, pp. 2185-2191. https://doi.org/ 10.1145/2695664.2695908

[4]M. Misbahuddin and B. Bindumadhava, "Design of a risk-based authentication system using machine learning techniques," in IEEE SmartWorld, 2017, pp. 149-200. https://doi.org/10.1109/UIC-ATC.2017.8397628

[5]H. Zhang, D. Singh, and X. Li, "Augmenting authentication with context-specific behavioral biometrics," in Proceedings of the Annual Hawaii International Conference on System Sciences, 2019, pp. 7282-7291. https://doi.org/10.24251/hicss.2019.875

[6]P. Arias-Cabarcos, C. Krupitzer, & C. Becker, "A survey on adaptive authentication", Acm Computing Surveys, vol. 52, no. 4, p. 1-30, 2019. https://doi.org/10.1145/3336117

[7]Ryu, R., Yeom, S., Herbert, D., & Dermoudy, J. (2023). A comprehensive survey of context-aware continuous implicit authentication in online learning environments. Ieee Access, 11, 24561-24573. https://doi.org/10.1109/access.2023.3253484

[8]S. Wiefling, T. Patil, M. Dürmuth, & L. Iacono, "Evaluation of risk-based re-authentication methods",, p. 280-294, 2020. https://doi.org/10.1007/978-3-030-58201-2_19

[9]S. Wiefling, M. Dürmuth, & L. Iacono, "More than just good passwords? a study on usability and security perceptions of risk-based authentication",, 2020. https://doi.org/10.1145/3427228.3427243

[10]R. Agrawal, "A study of touch dynamics biometrics authentication", Interantional Journal of Scientific Research in Engineering and Management, vol. 06, no. 05, 2022. https://doi.org/10.55041/ijsrem15812

[11]A. Buriro, B. Crispo, & M. Conti, "Answerauth: a bimodal behavioral biometric-based user authentication scheme for smartphones", Journal of Information Security and Applications, vol. 44, p. 89-103, 2019. https://doi.org/10.1016/j.jisa.2018.11.008

[12]R. Wang and D. Tao, "Context-aware implicit authentication of smartphone users based on multi-sensor behavior", Ieee Access, vol. 7, p. 119654-119667, 2019. https://doi.org/10.1109/access.2019.2936034

[13]A. Awwad, "An adaptive context-aware authentication system on smartphones using machine learning", International Journal of Safety and Security Engineering, vol. 13, no. 5, p. 903-915, 2023. https://doi.org/10.18280/ijsse.130514

[14]S. Fard, F. Gebali, & M. Mamun, "Using machine learning for dynamic authentication in telehealth: a tutorial", Sensors, vol. 22, no. 19, p. 7655, 2022. https://doi.org/10.3390/s22197655

[15]V. Zimmermann, P. Gerber, & A. Stöver, "That depends -- assessing user perceptions of authentication schemes across contexts of use",, 2022. https://doi.org/10.48550/arxiv.2209.13958

[16]S. Prange, L. Mecke, A. Nguyen, M. Khamis, & F. Alt, "Don't use fingerprint, it's raining!",, p. 1-5, 2020. https://doi.org/10.1145/3399715.3399823

[17]P. Arias-Cabarcos, C. Krupitzer, and C. Becker, "A survey on adaptive authentication," ACM Comput. Surv., vol. 52, no. 4, pp. 1–30, 2019. https://dl.acm.org/doi/10.1145/3336117

[18]K. A. Abu Bakar and G. R. Haron, "Adaptive authentication: Issues and challenges," in World Congress on Computer and Information Technology (WCCIT), IEEE, 2013, pp. 1–6. https://dl.acm.org/doi/10.1145/3582696

[19]G. Bendiab, N. Kolokotronis, S. Shiaeles, & S. Boucherkha, "Wip: a novel blockchain-based trust model for cloud identity management",, p. 724-729, 2018. https://doi.org/10.1109/dasc/picom/datacom/cyberscitec.2018.00126

[20]E. Ghazizadeh and B. Cusack, "Evaluation theory for characteristics of cloud identity trust framework",, 2019. https://doi.org/10.5772/intechopen.76338

[21]S. Hamdani, A. Khan, N. Iltaf, J. Bangash, Y. Bangash, & A. Khan, "Dynamic distributed trust management scheme for the internet of things", Turkish Journal

of Electrical Engineering & Computer Sciences, vol. 29, no. 2, p. 796-815, 2021. https://doi.org/10.3906/elk-2003-5

[22]S. Yanushkevich, W. Howells, K. Crockett, J. O'Shea, H. Oliveira, R. Guestet al., "Cognitive identity management: risks, trust and decisions using heterogeneous sources",, p. 33-42, 2019. https://doi.org/10.1109/cogmi48466.2019.00014

[23]S. Shao, "Master-slave multi-chain with risk assessment based access control model for zero trust network",, 2024. https://doi.org/10.21203/rs.3.rs-3869167/v1

[24]G. Bendiab, S. Shiaeles, & S. Boucherkha, "A new dynamic trust model for "on cloud" federated identity management",, p. 1-5, 2018. https://doi.org/10.1109/ntms.2018.8328673

[25]V. Varadharajan and S. Nepal, "Context-aware trust management system for iot applications with multiple domains",, p. 1138-1148, 2019. https://doi.org/10.1109/icdcs.2019.00116

[26]K. Awan, I. Din, A. Almogren, M. Guizani, A. Altameem, & S. Ullah, "Robusttrust – a pro-privacy robust distributed trust management mechanism for internet of things", Ieee Access, vol. 7, p. 62095-62106, 2019. https://doi.org/10.1109/access.2019.2916340

[27]]R. Aluvalu, K. Chennam, M. Jabbar, & S. Ahamed, "Risk aware access control model for trust based collaborative organizations in cloud", International Journal of Engineering & Technology, vol. 7, no. 4.6, p. 49, 2018. https://doi.org/10.14419/ijet.v7i4.6.20235

[28]T. hamme, D. Preuveneers, & W. Joosen, "Managing distributed trust relationships for multi-modal authentication", Journal of Information Security and Applications, vol. 40, p. 258-270, 2018. https://doi.org/10.1016/j.jisa.2018.01.003

[29]Q. Faxin, X. Tong, L. Yu, & Y. Wang, "Personalized project recommendations: using reinforcement learning", Eurasip Journal on Wireless Communications and Networking, vol. 2019, no. 1, 2019. https://doi.org/10.1186/s13638-019-1619-6

[30]S. Erat, H. Kitapçı, & K. Akçin, "Managerial perception and organizational identity: a comparative analysis", Sustainability, vol. 12, no. 6, p. 2278, 2020. https://doi.org/10.3390/su12062278

[31] D. Seng, B. Li, C. Lai, & J. Wang, "Adaptive learning user implicit trust behavior based on graph convolution network", Ieee Access, vol. 9, p. 108363-108372, 2021. https://doi.org/10.1109/access.2021.3100762

[32] K. Hasegawa, N. O'Brien, M. Prendergast, C. Ajah, A. Neves, & S. Ghafur, "Cybersecurity interventions in health care organizations in low- and middle-income countries: scoping review", Journal of Medical Internet Research, vol. 26, p. e47311, 2024. https://doi.org/10.2196/47311

[33] J. Kaur, S. Hasan, S. Orthi, M. Miah, M. Goffer, C. Barikdaret al., "Advanced cyber threats and cybersecurity innovation - strategic approaches and emerging solutions", Journal of Computer Science and Technology Studies, vol. 5, no. 3, p. 112-121, 2024. https://doi.org/10.32996/jcsts.2023.5.3.9

[34] G. Nguyen and M. Ha, "The role of user adaptation and trust in understanding continuance intention towards mobile shopping: an extended expectation-confirmation model", Cogent Business & Management, vol. 8, no. 1, 2021. https://doi.org/10.1080/23311975.2021.1980248

[35] J. Singh, "Zenith armor : advancing security with zero trust measures", Interantional Journal of Scientific Research in Engineering and Management, vol. 08, no. 04, p. 1-5, 2024. https://doi.org/10.55041/ijsrem31326

[36] F. Silva, "Evolving approaches in cybersecurity: metrics and human factors", International Seven Journal of Multidisciplinary, vol. 1, no. 2, 2024. https://doi.org/10.56238/isevmjv1n2-010

[37] W. Meng, K. Choo, S. Furnell, A. Vasilakos, & C. Probst, "Towards bayesian-based trust management for insider attacks in healthcare software-defined networks", Ieee Transactions on Network and Service Management, vol. 15, no. 2, p. 761-773, 2018. https://doi.org/10.1109/tnsm.2018.2815280

[38]J. H. Addae, X. Sun, D. Towey, et al., "Exploring user behavioral data for adaptive cybersecurity," User Modeling and User-Adapted Interaction, vol. 29, pp. 701-750, 2019. https://doi.org/10.1007/s11257-019-09236-5

[39]A. Hassan, B. Nuseibeh and L. Pasquale, "Engineering Adaptive Authentication," 2021 IEEE International Conference on Autonomic Computing and Self-Organizing Systems Companion (ACSOS-C), DC, USA, 2021, pp. 275-280, doi: 10.1109/ACSOS-C52956.2021.00068.

[40]Patricia Arias-Cabarcos, Christian Krupitzer, and Christian Becker. 2019. A Survey on Adaptive Authentication. ACM Comput. Surv. 52, 4, Article 80 (July 2020), 30 pages. https://doi.org/10.1145/3336117

[41]"Behavioral Analytics in Cybersecurity," CyberProtex, 2024. https://www.cyberprotex.com/blogs/september-08th-2024

[42]wso2, “GitHub - wso2/product-is: Welcome to the WSO2 Identity Server source code! For info on working with the WSO2 Identity Server repository and contributing code, click the link below.,” GitHub, Mar. 04, 2024. https://github.com/wso2/product-is (accessed Dec. 21, 2024).

[43]H. Fereidouni, Hafid, Abdelhakim Senhaji, D. Makrakis, and Y. Baseri, “F-RBA: A Federated Learning-based Framework for Risk-based Authentication,” arXiv.org, 2024. https://arxiv.org/abs/2412.12324v1 (accessed Apr. 20, 2025).