future internet

MDPI

Article

Blockchain–AI–Geolocation Integrated Architecture for Mobile Identity and OTP Verification

Gajasin Gamage Damith Sulochana [1] and Dilshan Indraraj De Silva [2,*]

[1] Faculty of Graduate Studies and Research, Sri Lanka Institute of Information Technology, Malabe 10115, Sri Lanka; ms24016766@my.sliit.lk
[2] Faculty of Computing, Department of Information Technology, Sri Lanka Institute of Information Technology, Malabe 10115, Sri Lanka
* Correspondence: dilshan.i@sliit.lk

Abstract

One-Time Passwords (OTPs) are a core component of multi-factor authentication in banking, e-commerce, and digital platforms. However, conventional delivery channels such as SMS and email are increasingly vulnerable to SIM-swap fraud, phishing, spoofing, and session hijacking. This study proposes an end-to-end mobile authentication architecture that integrates a permissioned Hyperledger Fabric blockchain for tamper-evident identity management, an AI-driven risk engine for behavioral and SIM-swap anomaly detection, Zero-Knowledge Proofs (ZKPs) for privacy-preserving verification, and geolocation-bound OTP validation for contextual assurance. Hyperledger Fabric is selected for its permissioned governance, configurable endorsement policies, and deterministic chaincode execution, which together support regulatory compliance and high throughput without the overhead of cryptocurrency. The system is implemented as a set of modular microservices that combine encrypted off-chain storage with on-chain hash references and smart-contract-enforced policies for geofencing and privacy protection. Experimental results show sub-0.5 s total verification latency (including ZKP overhead), approximately 850 transactions per second throughput under an OR-endorsement policy, and an F1-score of 0.88 for SIM-swap detection. Collectively, these findings demonstrate a scalable, privacy-centric, and interoperable solution that strengthens OTP-based authentication while preserving user confidentiality, operational transparency, and regulatory compliance across mobile network operators.

Keywords: OTP; SIM swap detection; geolocation authentication; artificial intelligence; blockchain identity

check for updates

Academic Editors: Eirini Eleni Tsiropoulou and Diala Naboulsi

Received: 10 October 2025
Revised: 18 November 2025
Accepted: 20 November 2025
Published: 23 November 2025

Citation: Sulochana, G.G.D.; De Silva, D.I. Blockchain–AI–Geolocation Integrated Architecture for Mobile Identity and OTP Verification. *Future Internet* **2025**, *17*, 534. https://doi.org/10.3390/fi17120534

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The rapid proliferation of mobile technologies has transformed digital identity verification and user authentication. Among the most widely used mechanisms is the One-Time Password (OTP), a dynamic, single-use code applied during login, transaction authorization, or other sensitive operations. OTPs form a central element of two-factor authentication (2FA) and are widely deployed in banking, e-government portals, e-commerce, and social media due to their simplicity and compatibility with mobile devices.

Despite this ubiquity, traditional OTP systems exhibit critical security weaknesses. SIM-swap attacks exploit operational gaps in Mobile Network Operator (MNO) procedures to reassign a victim's phone number, enabling attackers to intercept OTPs and gain unauthorized access to the victim's account. Phishing and automated OTP interception further

exacerbate this risk, as adversaries trick users into revealing OTPs in real time. These vulnerabilities arise largely from the reliance on centralized delivery channels, such as SMS and email, which are inherently insecure and lack contextual validation [1,2]. Recent analyses also show that SMS-based OTPs remain exposed to device-level malware and local attackers, emphasizing the fragility of centralized OTP delivery mechanisms [3].

The scale of resulting fraud is substantial. Reports from regulatory bodies and industry sources attribute billions of dollars in losses annually to authentication-related cybercrime, including SIM swapping and OTP exploitation. Account-takeover incidents leveraging SIM swaps or phishing have risen sharply, while OTP Phishing-as-a-Service (PaaS) platforms now enable industrialized, real-time interception of OTPs during live sessions. Tens of millions of users are affected each year, with cumulative global losses measured in tens of billions of dollars. These trends highlight the urgent need for context-aware, tamper-resistant, and privacy-preserving alternatives to conventional OTP delivery and verification.

Existing research has sought to strengthen OTP security using biometric factors, time-based cryptographic schemes, or enhanced token generation. However, most solutions address symptoms rather than underlying structural issues: centralized identity management [4], limited or absent real-time fraud detection, and poor interoperability among MNOs. Scalability and privacy constraints further limit deployment at telecom scale.

Addressing these challenges is both a technical and socio-economic imperative. As digital services continue to expand across borders, the financial and reputational damage caused by compromised authentication is growing. Ensuring secure yet flexible identity verification has therefore become a critical global priority.

This paper proposes an original, privacy-preserving framework for mobile identity verification that integrates blockchain technology, Artificial Intelligence (AI), and geolocation-based validation. The framework mitigates OTP vulnerabilities by combining decentralized ledger immutability, intelligent anomaly detection, and contextual verification. Unlike prior blockchain-based OTP models that rely on isolated smart contracts or centralized intermediaries, the proposed architecture emphasizes interoperability, data confidentiality, and real-time fraud intelligence.

The research integrates three core components: private blockchain, AI-driven anomaly detection, and geolocation-based OTP validation, combining them into a unified authentication framework. Hyperledger Fabric serves as a permissioned and tamper-evident ledger for decentralized identity management. Geo-hashing restricts OTP requests to legitimate physical regions, reducing risks associated with device spoofing and remote attacks. In parallel, a Machine Learning (ML) model continuously analyzes SIM usage behavior to detect SIM swaps and phishing-related anomalies in real time, even before OTP expiry.

The framework overcomes traditional silos by combining these technologies into a scalable and privacy-centric architecture. Fabric supports modular identity management through smart contracts and a hybrid data storage model, where only encrypted hashes of user and transaction data are stored on-chain and sensitive attributes remain securely off-chain. This balances transparency, integrity, and privacy. The AI-based anomaly detection module uses behavioral pattern analysis to flag abnormal SIM activity and synthetic anomalies linked to SIM swap attempts, while geolocation verification checks whether an OTP request originates from a user's typical location, applying geo-hash tolerance thresholds to reduce false positives.

This hybrid approach mitigates the weaknesses of conventional OTP mechanisms by reducing centralized dependencies [5] and enabling real-time fraud intelligence. It aligns with emerging decentralized-identity paradigms envisioned for Beyond-5G (B5G) ecosystems, which demand low-latency, high-security, and interoperable infrastructures.

The proposed architecture is adaptable across diverse mobile networks and service environments, emphasizing privacy protection and operational robustness.

Two telecom-specific methodologies are integrated within this framework:

- The MNO Decentralized Identity Network (TDIDN), which applies Decentralized Identifiers (DIDs) and blockchain-based consensus to enable secure, cross-network identity management without centralized repositories.
- The GSMA Mobile Identity Guidelines, which leverage mobile network attributes (e.g., SIM data, device IDs, and location metadata) for continuous risk assessment and adaptive authentication.

By combining these principles, the proposed model supports scalable, privacy-preserving, and context-aware mobile identity systems, establishing a foundation for verifiable and interoperable authentication in next-generation digital ecosystems.

This study introduces an end-to-end, privacy-preserving OTP authentication framework that integrates private blockchain, Zero-Knowledge Proofs (ZKPs), AI-driven behavioral analytics, and geolocation-aware validation. The architecture enhances both security and privacy while maintaining operational scalability. Its core innovations and validated contributions are as follows:

- ZKP–Fabric integration for privacy-preserving auditability: ZKP verification logic is embedded within the Hyperledger Fabric channel, enabling verifiers to confirm OTP transaction authenticity without accessing personally identifiable or geo-sensitive data. Only hashed proofs and credentials are stored on-chain, providing verifiable transparency while preserving privacy and maintaining an immutable, cryptographically sound audit trail.
- Configurable endorsement policies for scalable throughput: Fabric's modular endorsement mechanisms (e.g., OR/AND policies) are used to balance performance and security. Empirical evaluation shows that an optimized OR policy achieves approximately 850 transactions per second with ~220 ms commit latency, compared to 620 TPS and ~310 ms under a stricter AND policy, demonstrating that the framework maintains integrity while providing low-latency authentication for mobile environments.
- AI risk engine for adaptive SIM-swap and anomaly detection: A supervised anomaly detection model monitors SIM lifecycle patterns and network-level behaviors. The model substantially improves detection performance over a rule-based baseline, reducing missed SIM-swap events and enabling proactive defense against SIM swap and synthetic injection attacks within OTP workflows.
- Geolocation-bound OTP policy with contextual risk scoring: A GeoHash-based contextual scoring mechanism ensures that OTP requests originate from consistent spatial and behavioral contexts. Requests deviating beyond configured geo-tolerance thresholds trigger adaptive verification or rejection, preventing unauthorized OTP replay from distant or spoofed locations and adding a dynamic, context-aware security layer.
- Quantitative performance validation: Experimental results show that ZKP-enhanced verification introduces only ~160–190 ms of computational overhead while keeping total authentication latency under 0.5 s, which is acceptable for both banking and telecom use cases. This demonstrates that strong privacy guarantees via ZKPs can coexist with real-time operational performance.

Collectively, these contributions define a holistic, verifiable, and adaptive authentication framework that unifies blockchain auditability, AI-driven behavioral intelligence, and geolocation-context validation through cryptographic ZKP integration. This provides a level of combined assurance not achieved in prior studies.

2. Background Study

2.1. OTP-Based Authentication

OTP-based authentication is widely deployed across digital banking, e-commerce, and mobile applications as a lightweight mechanism for verifying user identity during sensitive operations. OTP codes are generated dynamically and remain valid for a single session or transaction, which helps mitigate replay attacks. Implementations typically follow either time-based or event-based models and are commonly delivered through SMS, email, or application-based generators.

Despite their broad adoption and ease of use, traditional OTP systems exhibit several structural weaknesses. Recent analyses of mobile authentication highlight persistent vulnerabilities in device-level token handling and multi-factor reliability, particularly when OTPs are transmitted via centralized delivery channels such as SMS or email [6]. High-impact attack vectors—including SIM swap fraud and phishing—exploit these centralized channels' lack of binding between the request context and the delivery endpoint [1,2], enabling adversaries to intercept or socially engineer access to OTPs. These issues underscore the need for contextual and behavior-aware OTP validation rather than relying solely on code possession.

To strengthen OTP security beyond conventional delivery mechanisms, research has explored more cryptographically robust designs. Innovations such as Merkle-tree–based one-time password schemes offer improved resistance to replay, interception, and leakage attacks by ensuring forward-security properties and minimizing exposure of token seeds [7]. In parallel, studies in mobile-centric fintech environments show a trend toward complementing OTPs with behavioral and contextual risk signals (e.g., device continuity, location consistency, or user interaction patterns) to improve robustness in high-risk transactions [8].

2.2. Blockchain-Based Identity Management

Conventional identity management architectures rely heavily on centralized authorities, which exposes sensitive user data to single points of failure, insider abuse, and large-scale data breaches [4,9]. These structural weaknesses have motivated significant interest in decentralized identity models, where trust is distributed and data integrity is verifiable without depending on a single administrative entity. Large-scale surveys of blockchain-based identity ecosystems confirm that decentralization improves auditability, reduces reliance on centralized trust anchors, and strengthens authentication resilience in multi-party environments [10].

Blockchain has therefore emerged as a compelling foundation for secure authentication and identity frameworks, offering tamper resistance, transparent traceability, and cryptographic guarantees of data integrity. Prior studies demonstrate blockchain's effectiveness in supporting robust multi-factor authentication and decentralized credential management, reducing risks associated with traditional identity silos [11–13].

Private, permissioned blockchain platforms such as Hyperledger Fabric are particularly suited for regulated environments that require controlled membership while maintaining verifiable, tamper-evident records. Fabric's modular architecture allows operators to configure endorsement, ordering, and access-control policies to achieve an appropriate balance between decentralization, performance, and governance. Experimental evaluations have shown that tuning endorsement policies significantly influences throughput. For example, moving from AND- to OR-based endorsement improves transaction throughput from 620 TPS to 850 TPS while maintaining acceptable latency.

Fabric was selected for this study due to three key advantages:

- Permissioned governance, supporting compliance-driven environments that require controlled identity issuance and auditable operations.
- Configurable trust and endorsement policies, allowing flexible security–performance trade-offs suited to MNO ecosystems.
- Deterministic smart-contract execution, which ensures consistent identity verification logic across consortium participants.

Compared with public blockchain networks, Hyperledger Fabric eliminates cryptocurrency-based transaction fees and provides built-in membership services with fine-grained, role-based access control. These capabilities make it particularly suitable for regulated identity verification and real-time OTP authentication. However, many blockchain-based identity systems continue to depend on centralized or semi-centralized off-chain storage for sensitive attributes, introducing residual trust dependencies that must be carefully managed [14–16]. Recent domain-specific blockchain authentication frameworks, including those designed for dynamic environments such as the Internet of Vehicles, demonstrate how decentralized identity records can strengthen trust and auditability, yet also highlight persistent challenges related to scalability and system overhead [17].

2.3. AI-Driven Anomaly Detection

AI-driven anomaly detection has become a fundamental component of modern authentication systems because of its ability to identify behavioral and contextual deviations that traditional OTP mechanisms cannot detect. By analyzing SIM lifecycle events, device fingerprints, network metadata, and temporal account activity, AI models can surface suspicious patterns—such as abrupt SIM reassignments, irregular OTP request timing, or atypical geospatial behavior—that often precede SIM-swap or account-takeover attacks [18]. These methods are particularly effective in telecom environments where fraud events are rare and datasets are highly imbalanced, enabling earlier detection of high-risk activity than rule-based approaches [19,20].

Recent advances leverage deep learning architectures—including recurrent neural networks, graph neural networks, and transformer-based sequence models—to capture temporal and relational dependencies that statistical techniques frequently overlook. Such models have demonstrated strong performance in detecting identity fraud, synthetic identities, and anomalous signaling patterns across heterogeneous telecom and IoT datasets [21,22]. In OTP workflows, these techniques enable continuous real-time monitoring of user behavior, device integrity, and contextual conditions, strengthening the system's ability to differentiate between legitimate and fraudulent requests [19,23].

To reduce false positives, AI-based systems increasingly incorporate contextual intelligence—such as subscriber history, device reputation, and transaction intent—resulting in more accurate risk evaluation with minimal friction for legitimate users. Hybrid strategies that combine unsupervised anomaly detection with supervised classification provide adaptability to emerging attack patterns while maintaining robustness against known threats [19,20,23]. In large-scale telecom networks, graph-based anomaly detection further enhances fraud detection by modeling relationships among devices, SIM identities, and network events, enabling the identification of coordinated attacks and malicious clustering activity [21,22].

Emerging research also highlights the potential of federated learning to support privacy-preserving anomaly detection by allowing models to be trained collaboratively across distributed datasets while retaining data locality. These approaches address scalability and adaptability challenges and provide promising avenues for secure, regulation-aligned fraud detection in mobile ecosystems [24].

Collectively, AI-driven anomaly detection significantly enhances the resilience of OTP-based authentication by proactively identifying high-risk events, modeling complex behavioral deviations in real time, and enabling dynamic adjustments to trust decisions before critical transactions are authorized.

2.4. Geolocation-Based Contextual Verification

Geolocation-based contextual verification augments OTP workflows by binding each authentication attempt to a plausible physical presence. In such systems, the verifier derives a coarse location representation (e.g., a geo-hash) from device-side signals and network metadata, and compares it against historical baselines and policy-defined geofences to decide whether to approve, challenge, or block the request. Surveys of indoor and outdoor positioning security show that carefully engineered location context can significantly strengthen access control while also imposing design requirements related to spoofing resistance and privacy protection [25].

For high-assurance scenarios requiring fine-grained spatial guarantees (e.g., branch offices, kiosks, or tightly controlled service areas), indoor localization can serve as an additional factor. LocAuth exemplifies this approach by using ambient Wi-Fi and Bluetooth characteristics, combined with supervised learning, to distinguish users within small indoor "trusted zones," achieving high classification accuracy in real deployments and making it suitable for OTP step-up authentication at regulated sites [26].

Because Global Navigation Satellite System (GNSS) readings alone are vulnerable to tampering (e.g., GPS spoofing), robust geolocation verifiers cross-check device-reported coordinates against independent signals. One practical strategy is multi-source consistency checking, which fuses cellular and Wi-Fi beacon observations with GNSS data to detect location forgery by flagging inconsistencies between crowd-sourced network proximity and claimed GPS positions [27]. Beyond radio-based cues, environmental fingerprints such as acoustic signatures can also help validate "you are here" claims: AcousticAuth, for example, authenticates stationary business devices based on venue-specific acoustic profiles, enabling zero-effort, location-anchored verification that is difficult to relay remotely [28].

Privacy remains a first-class requirement in geolocation-based authentication. Precise coordinates are often unnecessary to obtain meaningful risk signals. Research on privacy-preserving location-based services shows that systems can protect exact user positions while still preserving useful aggregate spatial utility (e.g., through distribution-aware obfuscation), enabling risk scoring and geofence checks without exposing point-level location data and aligning with regulatory expectations for minimal data processing in OTP contexts [29].

When verifiable and auditable proofs of presence are needed—such as in high-value operations that require strong non-repudiation—blockchain-based proof-of-location frameworks can deliver trustworthy spatio-temporal assertions. By using witness attestations and smart contracts, these systems generate tamper-evident location records that integrate with off-chain risk engines, offering an additional cryptographically verifiable layer of contextual assurance [30].

2.5. Common Security Threats

OTP-based authentication continues to face several well-established security threats that exploit weaknesses in traditional delivery channels and endpoint integrity. SIM swap fraud remains one of the most damaging attacks, allowing adversaries to redirect a victim's mobile number to a new SIM card and intercept SMS-delivered OTPs. Phishing attacks similarly undermine OTP workflows by coercing users into revealing authentication codes or credentials through deceptive social-engineering techniques. Attackers may also

employ device spoofing, using compromised or emulated devices to impersonate legitimate endpoints, or launch man-in-the-middle (MITM) attacks to intercept OTPs transmitted across insecure or compromised networks.

These ongoing threats highlight the need for a multi-layered security approach that integrates decentralized trust, intelligent threat detection, and contextual validation. The combined use of blockchain, AI-driven analytics, and geolocation-based verification directly mitigates weaknesses in legacy OTP systems by improving traceability, anomaly detection, and environmental consistency checks. This direction aligns with broader trends in emerging cybersecurity frameworks, which increasingly advocate distributed ledgers, machine-learning analytics, and contextual risk scoring to counter evolving identity-based attacks [11,12]. Moreover, architectural surveys of blockchain systems emphasize persistent challenges—such as risks related to ledger tampering, key-management exposure, and off-chain data linkage—that must be addressed to ensure secure and reliable decentralized authentication [31].

3. Literature Review

3.1. Traditional OTP Protocols and Their Limitations

OTP protocols have long served as a core component of two-factor authentication, providing a short-lived verification code sent via SMS, email, or application-based generators to supplement password-based login. Widely adopted standards such as HMAC-based OTP (HOTP) and Time-based OTP (TOTP), defined in IETF RFC 4226 and RFC 6238, are valued for their lightweight design, cross-platform compatibility, and minimal deployment overhead.

However, despite their maturity and ubiquity, traditional OTP mechanisms exhibit significant security limitations in modern threat environments. Because these systems depend on insecure communication channels such as SMS and email, OTPs remain vulnerable to interception through SIM-swap attacks, phishing schemes, malware, and account hijacking techniques that allow adversaries to reroute or capture verification codes.

A second limitation stems from the lack of contextual awareness. Legacy OTP protocols do not evaluate device identity, geolocation consistency, or behavioral patterns during authentication, making them incapable of distinguishing legitimate activity from anomalous or fraudulent attempts. This absence of contextual risk evaluation is particularly problematic in high-value environments such as digital banking and e-commerce, where attackers often exploit behavioral blind spots.

Newer implementations, including widely used app-based authenticators (e.g., Google Authenticator, Authy), provide improved usability and stronger cryptographic storage, but they remain susceptible to compromise if the device is lost, cloned, or infected with malware. Moreover, centralized OTP delivery platforms introduce single points of failure and scalability bottlenecks. Under network disruptions or targeted denial-of-service conditions, these centralized systems may delay OTP delivery or become temporarily unavailable, reducing reliability and undermining user trust.

Overall, while OTP protocols remain essential to contemporary authentication ecosystems, the literature consistently highlights their limitations in environments where attackers can exploit weak delivery channels, endpoint compromise, or absent contextual validation. These shortcomings motivate the shift toward multi-layered, context-aware, decentralized, and AI-assisted authentication frameworks.

3.2. SIM Swap Fraud: Nature, Impact, and Detection Strategies

SIM swap fraud is one of the most critical threats to OTP-based authentication, particularly in mobile-centric service environments. In this attack, adversaries exploit procedural

weaknesses within MNO workflows to unlawfully reassign a victim's mobile number to a SIM card under their control. Once the port-out or SIM replacement is approved, attackers gain full interception capability over SMS-delivered OTPs and related verification messages, enabling unauthorized access to financial accounts, e-commerce platforms, and digital identity systems [32].

The impact of SIM swap fraud is amplified by its stealth and sophistication. Victims are often unaware of compromise until fraudulent transactions have already occurred, and high-profile incidents in the financial and cryptocurrency sectors underscore the severity of the problem. Social engineering remains a dominant enabling factor: even with PIN-based safeguards or knowledge-based verification procedures, attackers routinely manipulate customer-service personnel to approve illicit SIM transfers [32].

To mitigate these risks, researchers have proposed behavioral and contextual detection mechanisms that operate before or during the authentication process. Roberts and Fisher [33] introduced a geolocation-based SIM swap detection method that flags discrepancies between a user's historical or registered location and the origin of an OTP request. Their findings demonstrate that geospatial inconsistencies can serve as early indicators of compromise, improving anomaly detection accuracy in high-risk transactions [33].

AI-driven detection models further enhance SIM swap defense by continuously monitoring SIM lifecycle behaviors—such as unusual changes in usage patterns, device transitions, or irregular access timing. These models identify deviations from established behavioral fingerprints and have shown promise in identifying SIM swaps even under imbalanced or sparse data conditions. To improve robustness, some studies introduce synthetic anomaly injection to augment training datasets and increase detection precision [34].

More recent work applies advanced behavioral analytics to capture subtle shifts in SIM activity. Features such as device relocation patterns, access-time anomalies, and SIM usage fingerprints have proven useful for recognizing emerging fraud behaviors in real time [35]. Despite these advances, operational challenges remain—especially the latency associated with cross-operator validation of number reassignments, which can delay decisive action. As a result, many existing approaches remain reactive, identifying fraud only after unauthorized access has occurred.

To address these limitations, the framework proposed in this study integrates AI-powered SIM swap detection with blockchain-backed identity management, geolocation verification, and contextual behavioral modeling. This unified, proactive architecture aims to prevent unauthorized OTP issuance at the access layer, delivering a resilient and adaptive defense against evolving SIM swap strategies.

3.3. Blockchain-Based Identity Management for Mobile Networks

Blockchain has become a reliable foundation for securing mobile identity systems due to its core properties of immutability, distributed consensus, and tamper resistance. Together, these features provide strong guarantees of trust, auditability, and transparent verification across participating MNOs [32]. In mobile authentication contexts, blockchain mitigates traditional weaknesses of centralized identity providers, including single points of failure, insider abuse, and interoperability limitations across operators. By recording identity assertions on an immutable ledger without exposing sensitive data, blockchain offers a privacy-preserving verification layer well-suited for distributed mobile networks.

Despite these advantages, many existing identity verification approaches still depend on centralized repositories to manage credentials. Such designs introduce high-value breach targets and synchronization bottlenecks when multiple MNOs participate in the ecosystem. Blockchain-based identity systems address this issue through cryptographically linked, decentralized validation workflows in which only authorized nodes participate, enabling

secure exchange of identity proofs without reliance on third-party intermediaries [32,36]. In this architecture, decentralization and consensus mechanisms work together to ensure data integrity, cross-network interoperability, and end-to-end security. These attributes are essential for defending against SIM swap, phishing, and replay attacks.

Several blockchain-based identity models have been proposed to meet the security and portability requirements of mobile networks. Xu et al. [37] introduced a redactable blockchain identity scheme using chameleon hashing to support General Data Protection Regulation (GDPR)–compliant mutability while maintaining tamper evidence. Kim, Suh, and Kwon [32] similarly emphasize selective mutability to provide privacy-compliant authentication with fine-grained user control over identity disclosures. Complementing these, Kumar, Singh, and Zhang [38] proposed a blockchain-based Mobile Number Portability (MNP) architecture enabling dynamic identity synchronization and cross-operator portability in multi-network environments.

To address scalability constraints, recent work adopts permissioned blockchain platforms such as Hyperledger Fabric. Its modular consensus mechanisms, private channels, and role-based access controls make it well-suited for high-volume identity workloads typical of telecom environments. Anderson and Lewis [39] demonstrated a microservices-oriented OTP infrastructure built on Fabric, showing how blockchain can improve scalability and interoperability for real-time authentication systems. Hybrid on-chain/off-chain storage models further reduce ledger overhead by storing only cryptographic hashes or pointers on-chain while maintaining sensitive identity data off-chain, supporting privacy-by-design and regulatory compliance [32,38].

Nonetheless, challenges remain: integrating blockchain with legacy authentication infrastructures, achieving high throughput under mobile latency constraints, and ensuring cross-chain compatibility and regulatory alignment continue to shape ongoing research directions [36]. These considerations reinforce the importance of designing scalable, interoperable, and privacy-preserving blockchain identity frameworks tailored to mobile ecosystems.

3.4. AI-Driven SIM Swap Detection and Fraud Analytics

The rise in SIM swap attacks presents a major threat to OTP-based authentication systems operating over mobile networks. By exploiting procedural vulnerabilities within MNO workflows, adversaries can fraudulently reassign a victim's mobile number to a new SIM card, enabling interception of OTPs and unauthorized account access [32]. Because traditional static authentication methods cannot evaluate dynamic behavioral patterns or detect abnormal SIM activity in real time, more adaptive and intelligent detection mechanisms are required.

Recent developments in ML and AI have shown strong potential for identifying fraud through behavioral and contextual analysis. Classical ML classifiers, including random forests, Support Vector Machines (SVMs), and decision trees, have been used to detect fine-grained anomalies in labeled fraud datasets [33,35]. However, fraud datasets remain highly imbalanced and difficult to curate, limiting the effectiveness of purely supervised learning approaches.

To address data scarcity, researchers have increasingly adopted unsupervised learning methods such as k-means clustering and isolation forests to identify anomalous activity in unlabeled data [35]. These models analyze a wide range of behavioral signals, such as device changes, SIM lifecycle transitions, OTP request frequency, and geolocation deviations, to identify high-risk patterns that may indicate SIM swap attempts. Synthetic anomaly injection techniques, which introduce artificial fraud examples during training, further

enhance model robustness by improving sensitivity to low-frequency but high-impact attack scenarios [34].

Behavioral analytics have also expanded into temporal modeling through sequence-based architectures such as Long Short-Term Memory (LSTM) networks and other Recurrent Neural Networks (RNNs). These models capture temporal dependencies in login sequences, device persistence, and access histories, improving the detection of coordinated or evolving attack behaviors that would be missed by static feature sets [35].

Building upon these advances, the framework proposed in this study integrates AI-driven SIM swap detection with contextual and geolocation intelligence to produce dynamic, real-time risk scores. By analyzing SIM usage patterns, mobility profiles, and network-level anomalies, the system proactively identifies suspicious OTP requests before authorization. All analytics operate on privacy-preserving data representations, ensuring alignment with regulatory requirements—including GDPR—while maintaining robust fraud detection capabilities.

3.5. Geolocation-Based Authentication and Contextual Validation

Geolocation-based authentication has emerged as a powerful complement to conventional OTP mechanisms by introducing a spatial dimension to identity verification. Traditional OTP systems cannot determine whether a request originates from an expected location, leaving users vulnerable to SIM swap, phishing, and device-spoofing attacks. By incorporating GPS coordinates, IP addresses, and cell-tower triangulation, geolocation-based systems add contextual intelligence that strengthens the verification process [33,35].

Unlike static authentication methods, geolocation-based approaches are adaptive and context-aware. They continuously evaluate user mobility patterns and can flag anomalous access attempts, such as logins from distant or high-risk regions, for additional verification. Roberts and Fisher [33] demonstrated that variations between typical SIM registration locations and the origin of OTP requests are strong indicators of potential fraud. Similarly, Nkwava et al. [40] showed that geofencing policies embedded in attendance and transaction platforms significantly reduced unauthorized access and location spoofing.

Geofencing itself operates by defining virtual boundaries (e.g., city-, country-, or organization-level zones) within which authentication attempts are considered legitimate. Banking and ATM authentication models proposed by Kumar, Singh, and Choudhary [41], as well as Alabdulatif et al. [42], automatically block OTP or transaction requests originating outside authorized regions. These contextual rules are particularly effective in high-risk environments, such as ATM withdrawals, cross-border payments, and international remittances, where location deviations often correlate with fraudulent intent.

Modern systems extend contextual validation beyond raw location data. By integrating supplementary signals such as motion indicators, environmental cues, Bluetooth beacon data, and Wi-Fi triangulation, these systems strengthen their resistance to spoofing. Adams [43] introduced GeoAuth, a browser-based framework that fuses real-time geolocation with contextual session indicators to differentiate legitimate user activity from anomalies associated with phishing or credential theft.

Despite their advantages, geolocation-based authentication methods face operational and privacy-related challenges. GPS spoofing, VPN masking, and poor signal conditions can degrade accuracy, while continuous location tracking raises ethical and regulatory concerns. These issues can be mitigated through consent-based data collection, anonymization, and minimal disclosure practices aligned with privacy regulations.

In the proposed framework, geolocation operates alongside AI-driven behavioral analytics and blockchain-based identity management to provide layered, context-rich OTP validation. OTPs are issued only when the user's current geolocation aligns with

reference patterns stored on-chain or predicted by the AI model. This integration blocks unauthorized remote access, supports real-time anomaly detection, and ensures compliance with regulatory frameworks such as the GDPR.

3.6. Integrated Framework: Combining Blockchain, AI, and Geolocation

To combat increasingly sophisticated authentication threats—including SIM swap attacks, phishing, and device spoofing—this study proposes an integrated architecture that combines blockchain, AI-driven analytics, and geolocation-based contextual validation. Each component addresses a different dimension of risk, and together they form a cohesive, verifiable, and adaptive authentication ecosystem.

The framework employs a permissioned blockchain built on Hyperledger Fabric to provide decentralized identity management, smart-contract–based policy enforcement, and tamper-resistant auditability. Identity information is stored using a hybrid model in which cryptographic hashes reside on-chain for integrity and non-repudiation, while sensitive identity attributes remain off-chain to preserve confidentiality and meet privacy regulations such as the GDPR. This separation ensures that verification processes benefit from blockchain immutability without disclosing personal data.

Within this architecture, AI-driven anomaly detection continuously evaluates behavioral, device-level, and SIM-related indicators to identify fraud patterns before OTP issuance. Concurrently, geolocation intelligence verifies whether an authentication request originates from a legitimate and contextually plausible location. Together, these modules provide real-time, multidimensional insight into user authenticity, enabling proactive and privacy-preserving decision-making.

Integrated with the blockchain layer, the AI-powered fraud analytics engine applies both supervised and unsupervised learning methods to detect irregularities in SIM lifecycle events, OTP request timing, device identity shifts, and mobility patterns. Synthetic anomaly injection is incorporated to compensate for the scarcity of labeled fraud data, improving model robustness and generalization when confronting rare but high-impact attack scenarios [34,35].

The geolocation validation module acts as a contextual filter within the authentication pipeline. It compares the user's real-time location with historical access patterns stored on-chain or modeled by the AI engine. OTPs are issued only when the observed location aligns with trusted geospatial boundaries. This module employs geofencing rules and dynamic risk scoring to strengthen defenses against location spoofing, remote account takeover attempts, and geographically inconsistent OTP requests [33,41,44].

The integrated workflow proceeds as follows:

- A user initiates an OTP request via mobile or web.
- The AI engine analyzes the request using behavioral history, device signatures, and timing anomalies.
- Simultaneously, the geolocation module verifies whether the request originates from an authorized region.
- A smart contract evaluates blockchain-stored identity anchors and contextual signals from both modules.
- If all policies are satisfied, the OTP is issued; otherwise, the system blocks the request or escalates it for manual review.

By decentralizing identity data, applying predictive fraud analytics, and incorporating spatial and contextual intelligence, the framework transitions authentication from a reactive, credential-centric process to a real-time, adaptive security mechanism. Its modular design enhances scalability, auditability, and resilience—making it particularly suitable for high-

risk sectors such as digital banking, fintech, and mobile commerce, where autonomous and verifiable authentication is essential for maintaining trust and regulatory compliance.

To contextualize the integrated framework within the broader body of related work, Table 1 provides a consolidated summary of the key papers reviewed in this study, highlighting their focus areas, evaluation methods, addressed threats, principal findings, and observed limitations.

Table 1. Comparative overview of research studies on OTP authentication, SIM swap defense, geolocation-based security, and blockchain identity frameworks.

Ref.	Focus Area	Techniques	Evaluation Methods and Metrics	Security Threats Addressed	Key Findings	Limitations
[32]	SIM swap mitigation	Survey of cases; procedural and technical controls	Narrative review; no empirical evaluation reported	SIM swap, social engineering	SIM lifecycle monitoring and multi-party verification reduce SIM swap risk	Lack of real-world datasets; regional bias possible
[33]	Geolocation auth	Predefined geographic zones; policy-based blocking	Prototype; qualitative analysis	Account takeover, location spoofing	Out-of-zone attempts are strong risk indicators	Usability trade-offs; false positives when traveling
[34]	Geolocation auth	Server–Assisted Matching of user location with policy	Security analysis; prototype	MITM, phishing	Combining network and GPS signals strengthens auth	Accuracy variations across environments
[35]	SIM swap mitigation, Geolocation auth	Use SIM Toolkit + user location as OTP gate	Design and feasibility; limited measurements	SIM swap	Location checks before OTP issuance block SIM swap interception	Coverage and privacy constraints
[36]	Blockchain identity, MNP	Blockchain-based subscription and porting; off-chain storage	Prototype; latency/throughput tests	Centralized MNP, tampering	DLT removes single point of failure; improves auditability	Requires governance; tested only in simulated environment
[37]	Blockchain identity	Redactable blockchain enabling GDPR-aligned erasure	Security proofs; prototype	Privacy, insider threats	Selective mutability with verifiable integrity	Governance of redaction keys
[38]	OTP security	TLS-based seed exchange; encrypted keystore	Security model; performance tests	Seed theft, SMS interception	Secure seed provisioning mitigates SIM/SMS risks	Device compromise risk remains
[39]	OTP infrastructure	Microservices; scalable OTP issuance/verification	Reference implementation; throughput tests	DoS, single-point failure	Decoupling improves resilience and scalability	Security depends on surrounding controls
[40]	Blockchain identity, MNP	Distributed ledger for cross-operator profile/MNP tracking	Architecture and simulation	Port-out fraud, insider tampering	DLT improves auditability and portability trust	Interoperability; performance overhead
[41]	Geolocation auth	JWT auth, dynamic QR, geofencing, IMEI checks	Prototype; MERN stack; case testing	Location spoofing, unauthorized access	Geo-bound enforcement prevents off-site submissions	GPS/network dependency
[42]	ATM auth, Geolocation auth	ATM transactions allowed only within user's geofence	Prototype; scenario tests	Card cloning, shoulder surfing	Geo-bound MFA reduces off-site fraud	User mobility; network availability
[43]	Geolocation auth	Context-aware browser; real-time device location and context checks	Prototype-based security analysis	Phishing, session hijack	Binding auth to physical context reduces fraudulent sessions	Location spoofing risk; GPS/network bias
[44]	Geolocation auth, OTP security	Predict next user location to pre-validate OTP requests	Algorithmic evaluation (latency/accuracy)	Phishing, remote takeover	Low-latency prediction enables risk-based OTP delivery	Model drift; outlier behavior
[45]	OTP security	Survey of HOTP/TOTP and variants	Comparative analysis	Channel interception, replay	TOTP widely adopted but vulnerable in SMS delivery	No new defenses implemented
[46]	Fraud prevention (fintech)	Best-practice controls; AI-driven monitoring	Industry survey	Account takeover, bot abuse	Layered controls outperform single-factor defenses	Generic guidance; not telco-specific
[47]	OTP security, Biometrics	Derive OTP seed from fingerprint features	Accuracy/security analysis	Credential theft	Biometric-bound OTP raises entropy	Biometric privacy; spoofing defenses

3.7. Comparative Synthesis of Related Work

The reviewed literature collectively demonstrates a broad spectrum of approaches aimed at strengthening OTP-based authentication through cryptographic, behavioral, contextual, and decentralized mechanisms. Across these works, three dominant enhancement directions emerge:

- Protocol-level improvements to traditional HOTP/TOTP schemes;
- Architectural and infrastructural redesigns using blockchain or microservices;
- Context-aware and biometric extensions designed to mitigate emerging threats such as SIM swapping, phishing, and device compromise.

Early studies evaluating OTP protocols in isolation emphasize their resilience against phishing and MITM attacks but note the absence of built-in privacy protections and contextual validation mechanisms [45]. More advanced protocol-level enhancements rely on TLS-based OTP exchange and encrypted offline keystores, which demonstrate reduced latency and improved confidentiality in prototype deployments [46]. However, these approaches still depend on endpoint integrity and therefore remain susceptible to device-level compromise.

Biometric-bound OTP generation, such as fingerprint-derived seed creation, offers a stronger defense against impersonation attacks but exhibits inherent deployment constraints. Its reliance on specialized hardware restricts scalability, particularly in heterogeneous mobile environments or regions with limited biometric infrastructure [47]. Similarly, microservice-based OTP delivery frameworks provide improved throughput and fault isolation, yet they generally lack native privacy-preserving or cross-operator verification capabilities, limiting their applicability in multi-network settings [39].

A more recent research direction explores the integration of blockchain for decentralized identity management and OTP verification. These systems mitigate the risks associated with centralized repositories, including single points of failure, credential leakage, and insider tampering. They achieve this by distributing identity proofs across tamper-resistant ledgers. Several works incorporate off-chain storage mechanisms (e.g., IPFS) to support privacy-preserving data management, pseudonymous verification, and scalable record retrieval. While promising, these blockchain-enabled frameworks remain at early developmental stages and require comprehensive performance benchmarking under real mobile network conditions.

Parallel advancements embed geolocation and contextual intelligence directly into OTP workflows. Such studies demonstrate that geolocation constraints can effectively block spoofed or remote attacks targeting high-risk transactions. These constraints are implemented through geofencing, location-policy enforcement, or real-time location matching. A subset of these works combines blockchain with geolocation-based triggers, leveraging smart contracts to automate contextual validation when detecting SIM swap, device spoofing, or anomalous mobility patterns. Although these designs strengthen OTP integrity, they still face challenges related to GPS spoofing resistance, cross-border mobility, and environmental variability.

AI-driven OTP authentication models contribute an additional defense layer by performing anomaly detection using behavioral, temporal, and mobility-based features. These frameworks show particular value in detecting low-frequency but high-impact threats such as SIM swap attempts and coordinated account takeovers. Nevertheless, many remain confined to lab-scale evaluations due to limited access to real-world telecom datasets and the difficulty of validating detection accuracy under dynamic network conditions.

Taken together, these prior works illuminate significant progress in decentralization, privacy preservation, behavioral monitoring, and contextual authentication. However, most existing solutions focus on isolated components, including blockchain-only, AI-only, or geolocation-only defenses. They rarely integrate these elements into a unified, multi-layered framework. Furthermore, cross-operator interoperability, support for MNP environments, and compliance with international data-protection requirements remain underexplored. The present study builds upon these research trajectories by synthesizing blockchain, AI-

based anomaly analytics, and geolocation validation into a cohesive architecture designed for real-world deployment across heterogeneous mobile networks.

4. Methodology

This section outlines the system's architecture, processes, datasets, and evaluation protocol used to design and validate a privacy-preserving OTP authentication framework. The proposed model integrates a permissioned blockchain (Hyperledger Fabric) for tamper-evident identity auditing, an AI-driven risk engine for SIM-swap and anomaly detection, and geolocation-based contextual validation for policy-gated OTP issuance.

The framework explicitly supports mobile identity management across multiple MNOs via a consortium ledger that anchors pseudonymous identity proofs and SIM-lifecycle events (e.g., swaps or ports). It standardizes inter-operator APIs and enforces cross-operator policies through smart contracts.

This section defines the threat model and security objectives, explains the hybrid on/off-chain data-protection architecture, and details AI and geolocation validation controls. It further describes model features, training, and thresholds, outlines microservice interfaces for relying parties and MNO backends, and presents the evaluation approach using Caliper and JMeter within a reproducible, containerized testbed.

4.1. Architectural Overview

Figure 1 presents the high-level system architecture, illustrating the layered components and their interactions among relying parties and MNOs. The framework issues and validates OTPs only when identity, behavioral, and contextual criteria satisfy predefined trust thresholds. It is composed of five primary layers, each responsible for a distinct functional role:

- Identity and audit layer: Implemented using a permissioned Hyperledger Fabric network, this layer maintains tamper-evident records of pseudonymous identity proofs, policy decisions, and OTP-related events. Only salted hashes and coarse digests are written on-chain; sensitive attributes remain off-chain to support data-minimization requirements.
- Risk and intelligence layer: This layer evaluates OTP requests by generating real-time anomaly scores based on SIM-lifecycle events, mobile number portability (MNP) data, device-continuity indicators, and geo-behavioral patterns. Results are supplied to the policy engine for adaptive decision-making.
- Context layer: Geolocation validation is performed using geo-hashing and corroborated by network-based and device-based signals (e.g., GPS, cellular triangulation). A compact "geo-digest" is produced to summarize spatial consistency, movement plausibility, and agreement across data sources.
- Access layer: This layer exposes API-based microservices that support user registration, OTP request and verification, risk evaluation, and policy enforcement. It integrates operator event streams and provides OAuth2/OIDC endpoints for relying parties such as financial institutions and government services.
- Client layer: A lightweight mobile/web Software Development Kit (SDK) securely binds user devices to identity anchors through protected keystores and attestation mechanisms. It optionally supports app-based OTP generation when SMS delivery is restricted by policy.

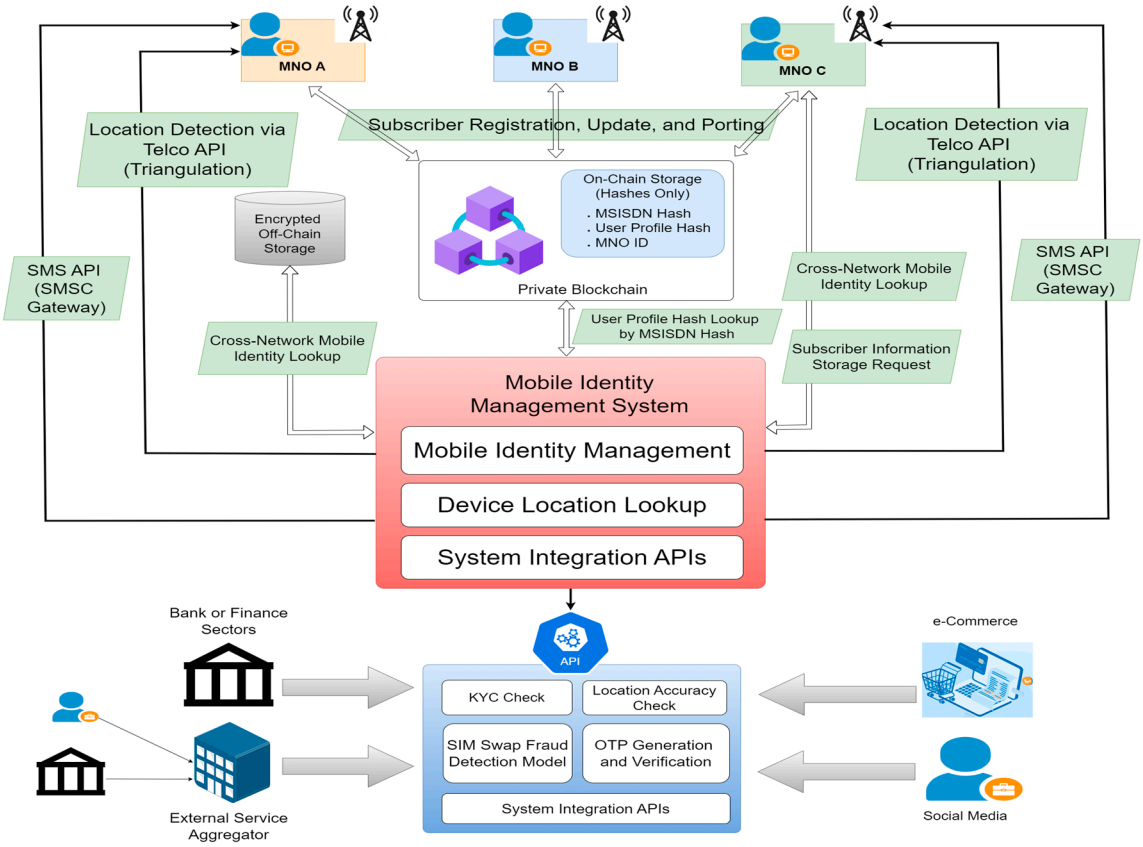

Figure 1. Overview of proposed architecture.

A key design goal is enabling cross-operator identity management without revealing user-identifiable attributes. Operators publish verifiable attestations of SIM-lifecycle and porting events to the consortium ledger, and smart contracts standardize schemas and policies across participants.

End-to-end processing follows a deterministic workflow: when a user requests an OTP, the policy engine evaluates trust parameters, records a tamper-evident receipt on the ledger, and authorizes or rejects OTP issuance based on risk conditions. The response to relying parties includes structured indicators such as KYC match summaries, SIM-swap history, and geolocation confidence scores. These indicators allow relying parties to apply their own acceptance thresholds or trigger additional verification. Security and privacy controls include the following:

- On-chain minimalism: Commit only non-sensitive proofs and digests.
- Mutual TLS and attestation: Authenticate all inter-service communications.
- Client integrity checks: Implement jailbreak/root detection and runtime integrity verification.
- Network-edge protections: Apply rate-limiting and application-layer filtering.

This layered architecture supports auditability through the blockchain layer, anomaly detection through the risk engine, and contextual validation through the geolocation layer. Together, they provide a structured foundation for privacy-preserving, cross-operator OTP authentication. Figure 2 depicts the OTP verification flow, detailing interactions among the client application, identity management system, MNO backend, fraud-detection engine, and blockchain verifier.

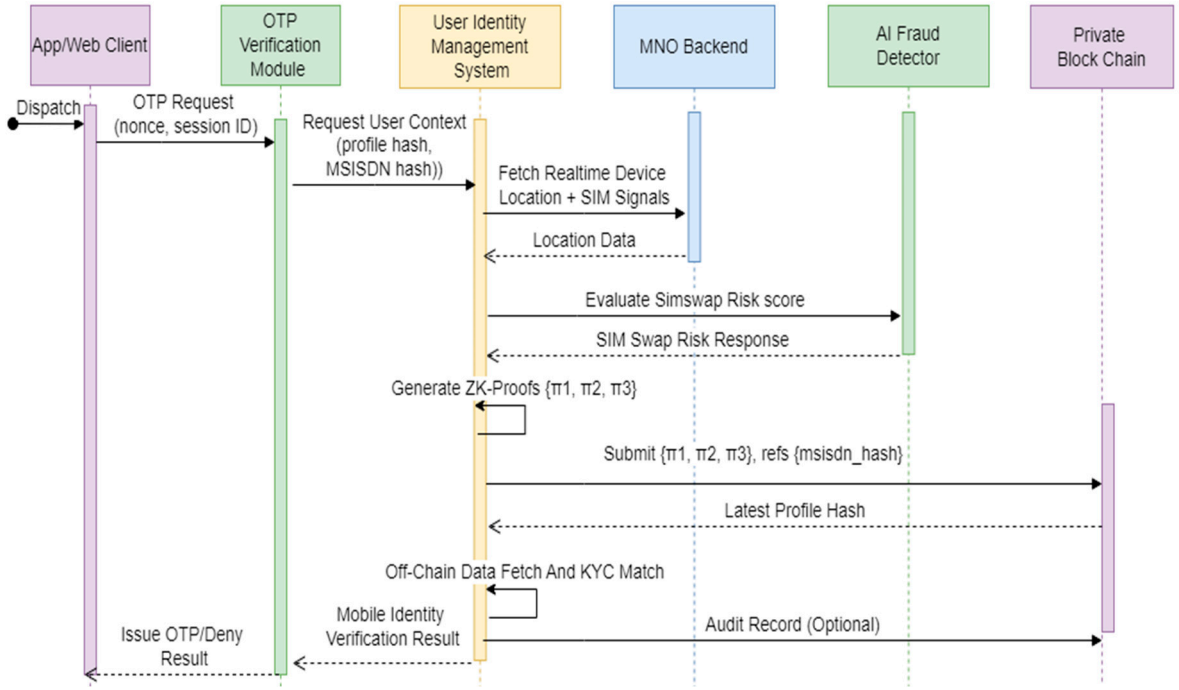

Figure 2. OTP verification sequence across the App/Web Client, OTP Verification Module, User Identity Management System, MNO backend, AI Fraud Detector, and Private Blockchain. Solid arrows represent primary data and control flows, while dashed arrows indicate response messages.

4.2. Zero-Knowledge Proof Mechanism for Privacy-Preserving Verification

The framework employs a ZKP mechanism to verify identity attributes, location consistency, and SIM-swap risk without disclosing raw personal data. The design integrates a zk-SNARK construction, enabling efficient, threshold-based assertions for key validation steps.

- Cryptographic construction: The system uses a zk-SNARK protocol (e.g., Groth16) deployed over a pairing-friendly elliptic curve. Poseidon or Pedersen hash functions are embedded in the proving circuit to support efficient commitments and minimize computational overhead.
- Statements proven: Three threshold-based statements are encoded in the circuit:
  1. S1: Identity consistency
     KYC_match $\geq \theta$, where similarity is computed over tokenized KYC attributes associated with the user's profile hash.
  2. S2: Location consistency
     geohash_match $\geq \theta$, confirming that the user's current device location is within an expected geofence or cell cluster.
  3. S3: SIM-swap risk compliance
     risk_score $\leq \tau$, where the AI-generated risk value is quantized to reveal only its threshold compliance.
- Prover and verifier roles:
  1. Prover: The User Identity Management System, operating off-chain, generates proof bundles $\{\pi_1, \pi_2, \pi_3\}$ using private inputs.
  2. Verifier: Hyperledger Fabric chaincode verifies proofs on-chain, producing deterministic and auditable outcomes.
- Public and private inputs:

1. Public inputs: Profile hash, MSISDN hash, validation thresholds (θ, τ), and the zk-SNARK verification key.
2. Private witnesses: Tokenized KYC vectors, device-location snapshots, feature vectors used by the fraud-detection model, and the raw risk score.

- Proof outputs: Consist of the ZK proof set and associated verification verdicts, with only hashed audit logs stored on-chain.
- Architectural integration: Figure 3 illustrates how the User Mobile Identity Management System (UMIMS) integrates with external MNOs, the OTP module, and Hyperledger Fabric. The sequence encompasses KYC verification, SIM-swap risk assessment, ZKP generation, and privacy-preserving OTP issuance across multiple trust domains.
- Privacy and limitations: The architecture ensures separation between public cryptographic anchors and private witness data. Only hashed identifiers (e.g., profile_hash and MSISDN hash) are committed to the ledger. However, several limitations persist:
  1. Schemes requiring trusted setup (e.g., Groth16) depend on secure parameter generation.
  2. Proof generation introduces non-trivial computational overhead on the prover.
  3. Circuits must be defined in advance, constraining flexibility when authentication logic or risk thresholds evolve.
  4. Current assumptions do not fully account for side-channel leakage during proof generation.
  5. The security of private data remains dependent on the integrity of the prover, which processes raw KYC and location information before producing proof.

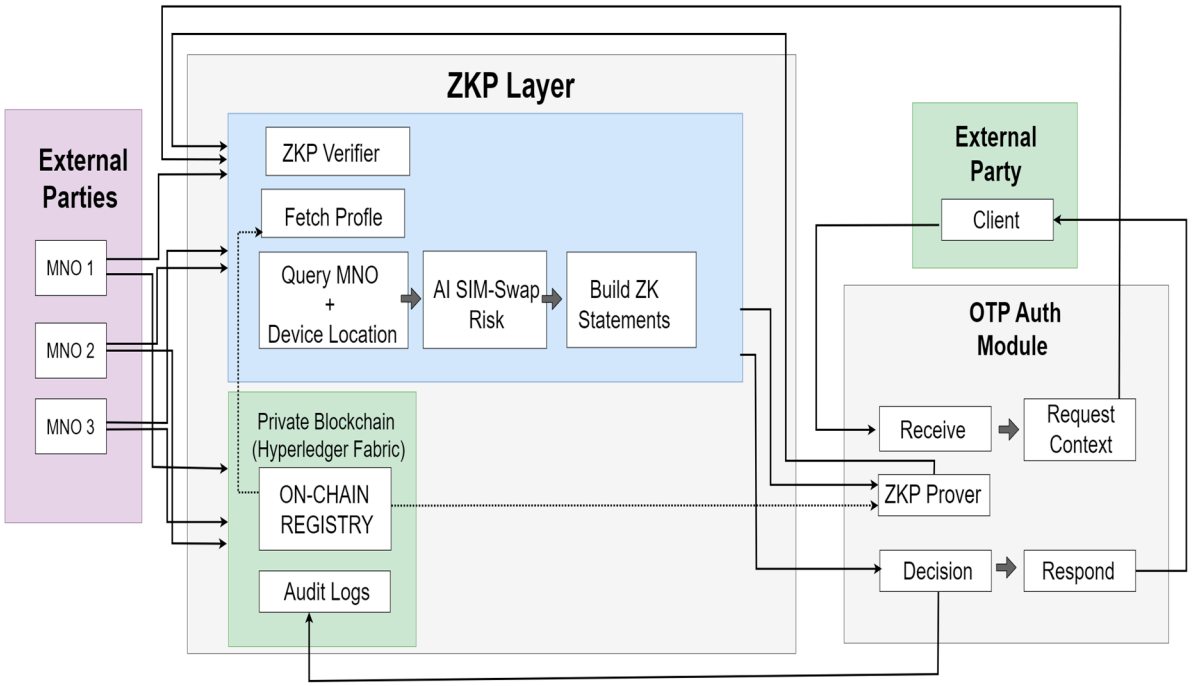

Figure 3. Proposed ZKP-enabled OTP authentication framework. Solid arrows represent primary data flows between components, while dashed arrows indicate data-query operations used during verification.

Future work will explore transparent or universal proof systems (e.g., Plonk, Halo2) and post-quantum–resistant ZKPs to enhance scalability and long-term security.

4.3. Verification Workflow, Algorithms, and Validation Examples

The proposed approach is formally validated through zk-SNARK proof verification, where each proof π_i is checked on-chain under the constraints KYC $\geq \theta$, geohash $\geq \theta$, and risk $\leq \tau$. Verification is performed by Hyperledger Fabric chaincode using elliptic-curve pairing operations, ensuring deterministic correctness and reproducibility.

4.3.1. Algorithm 1: OTP Generation with ZKP-Based Attribute Validation

Algorithm 1 defines the procedure for generating an OTP while simultaneously verifying KYC similarity, geolocation consistency, and SIM-swap risk. The algorithm runs within the User Identity Management System, which acts as the ZKP prover.

Algorithm 1: OTP_AND_PROOFS(user_id, session_nonce)

```
Inputs:
  user_id, session_nonce
  θ_KYC  // KYC similarity threshold
  θ_LOC  // location match threshold
  τ_SIM  // maximum acceptable SIM swap risk

Outputs:
  otp_token
  proofs = {π1, π2, π3}
  public_refs = {profile_hash, msisdn_hash}
-----------------------------------------------
Initialization and Data Retrieval
-----------------------------------------------
1: profile_hash  ← get_profile_hash(user_id)
2: msisdn_hash ← get_msisdn_hash(user_id)
3: kyc_vec_user ← fetch_user_input_kyc(user_id)          // tokenized
4: kyc_vec_ref  ← fetch_ref_kyc(profile_hash)            // tokenized
5: KYC_match  ← similarity(kyc_vec_user, kyc_vec_ref) // e.g., Jaccard or Cosine
6: loc_now       ← MNO.get_device_location(user_id)     // cell/triang./GPS snapshot
7: loc_ref       ← get_allowed_geofence(profile_hash)   // polygon or Geohash prefix
8: LOC_match  ← geofence_score(loc_now, loc_ref)       // range: [0, 1]
9: features ← build_features(user_id, msisdn_hash, loc_now,
device_fingerprint, sim_events)
10: risk_score ← SIM_SWAP_MODEL.predict(features)
-------------------------------------------------------
Construct ZK Statements and Witnesses
-------------------------------------------------------
11: stmt1 ← (KYC_match ≥ θ_KYC)
12: stmt2 ← (LOC_match ≥ θ_LOC)
13: stmt3 ← (risk_score ≤ τ_SIM)
14: witness ← { commit(kyc_vec_user), commit(kyc_vec_ref), commit(loc_now),
    commit(features), risk_score}
15: public_inputs ← { msisdn_hash, θ_KYC, θ_LOC, τ_SIM}
-------------------------------------------------------
Generate ZK Proofs and OTP
-------------------------------------------------------
16: {π1, π2, π3} ← zkSNARK.Prove({stmt1, stmt2, stmt3}, witness, public_inputs)
```

Algorithm 1: OTP_AND_PROOFS(user_id, session_nonce)

17: otp_token ← HOTP_or_TOTP_Derive(session_nonce,
key = KDF(profile_hash ‖ device_binding))
18: return otp_token, {π1, π2, π3}, {profile_hash, msisdn_hash}

4.3.2. Algorithm 2: SIM-Swap Fraud Detection Model

Algorithm 2 defines the SIM-swap risk model used to produce the risk score incorporated into ZKP Statement S3. The model evaluates a combination of network, device, and behavioral features. The algorithm is structured into the following:

- Training phase: Data cleaning, feature derivation, class-imbalance mitigation, model training (e.g., gradient-boosted trees), probability calibration, and threshold selection.
- Inference phase: Real-time feature construction, risk-score computation, and a binary decision outcome based on threshold τ.
- ZKP binding: The resulting risk score is quantized and provided as a private witness in the ZKP circuit, enabling an on-chain assertion that risk $\leq \tau$ without revealing the exact score.

Algorithm 2: SIM Swap Fraud Detection Model

```
Inputs (per record):
  sim_swap_time_gap_minutes              // minutes since last SIM change
  device_change_flag                     // {0, 1}
  sim_type_change_flag                   // {0, 1}
  imsi_change_flag                       // {0, 1}
  iccid_change_flag                      // {0, 1}
  otp_and_sim_change_geo_hash_length     // shared Geohash prefix length
  sim_swap_flag                          // {0, 1}  ground-truth label

Outputs:
  risk_score ∈ [0, 1]
  decision ∈ {ACCEPT, DENY}              // threshold τ for ZKP Statement S3
--------------------------------------
Training Phase (offline)
--------------------------------------
D ← collect labeled records with all input features
Clean and preprocess:
    Convert boolean indicators to {0, 1}
    Impute missing numeric values (median) for time_gap and geo_hash_length
Feature scaling:
    s_time  ← MinMaxScale(sim_swap_time_gap_minutes)
    s_geo   ← MinMaxScale(otp_and_sim_change_geo_hash_length)
Derived features:
    time_urgency_score ← 1 − s_time
    identity_shift_score ← device_change_flag
                           + sim_type_change_flag
                           + imsi_change_flag
                           + iccid_change_flag
    geo_risk_score       ← 1 − s_geo
```

Algorithm 2: SIM Swap Fraud Detection Model

```
5: Construct feature vector:
     X ← [ time_urgency_score,
           identity_shift_score,
           geo_risk_score,
           sim_swap_time_gap_minutes,
           otp_and_sim_change_geo_hash_length,
           device_change_flag,
           sim_type_change_flag,
           imsi_change_flag,
           iccid_change_flag ]
     y ← sim_swap_flag
6: Handle class imbalance (if present):
     Apply class_weight = "balanced"
     OR apply SMOTE on minority class
7: Train classifier:
     model ← GradientBoostedTrees (or XGBoost/Logistic Regression with interactions)
8: Probability calibration:
     calibrated_model ← Isotonic (or Platt) calibration on validation set
9: Select operating threshold τ:
     Scan τ ∈ [0, 1] to optimize target metric
     (e.g., minimize FN at fixed FP, or maximize F1 or Youden index)
     Persist τ, scaler parameters, and calibrated_model
-----------------------------------
Inference Phase (online)
-----------------------------------
10: Build feature vector X_now using the same steps as in 2–5
11: risk_score ← calibrated_model.predict_proba(X_now)[SIM_SWAP]
12: if risk_score > τ then
       decision ← DENY
    else
       decision ← ACCEPT
    end if
13: return risk_score, decision
-----------------------------------------
ZKP Binding (Statement S3)
-----------------------------------------
14: Provide risk_score (quantized/bucketed) as private witness to ZKP circuit
15: Publicly assert Statement S3: risk_score ≤ τ
16: If the ZKP verifies on-chain, Statement S3 passes;
     otherwise, OTP request is rejected
```

4.3.3. Validation Examples

Figures 4–6 illustrate how the verification pipeline handles inconsistent or anomalous identity, location, and SIM-swap indicators:

- KYC mismatch (Figure 4): During onboarding, a KYC match score of 27.27%, which is below the configured threshold of 60%, indicates divergence between the submitted user data and the MNO-verified subscriber records.

- Location inconsistency (Figure 5): OTP requests originating from locations significantly distant from the transaction source are rejected when they fall outside the permitted location-accuracy threshold. Threshold values are configurable by relying parties.
- SIM-swap risk (Figure 6): Behavioral anomalies, such as recent IMSI/ICCID changes, device inconsistencies, or unusual geolocation patterns, may produce elevated risk scores (e.g., 50% against a 25% threshold), triggering secondary verification during registration or authentication.

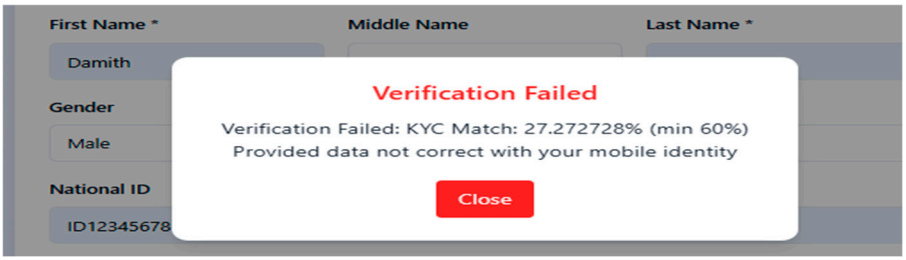

Figure 4. Failed KYC verification due to a mismatch between the user-submitted data and the MNO-verified subscriber identity.

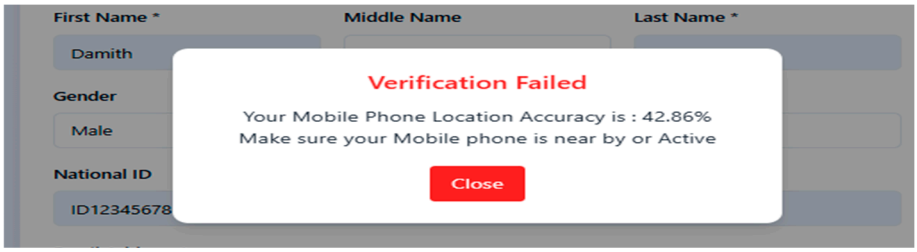

Figure 5. Failed geolocation consistency check indicating insufficient proximity between the device location and the expected transaction source.

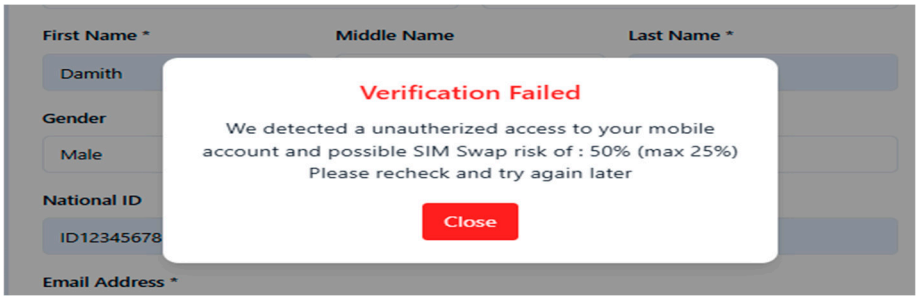

Figure 6. SIM-swap risk alert triggered by anomalous device, SIM, or geolocation indicators exceeding the configured risk threshold.

Together, these examples show how the framework integrates KYC verification, geolocation checks, and SIM-swap risk assessment into a unified decision process. Verification layers can be selectively activated through policy controls, enabling adaptive and context-aware authentication consistent with privacy-preserving design principles.

4.4. Threat Model and Security Objectives

4.4.1. Threat Model

The proposed authentication framework assumes a hostile operating environment in which adversaries may possess varying levels of access, technical capability, and opportunity. The threat model characterizes system assets, trust boundaries, and adversarial capabilities, followed by the resulting threat categories relevant to OTP-based authentication.

- System assets—the primary assets requiring protection include the following:
  1. User identity attributes (e.g., hashed identifiers, device fingerprints, mobility profiles)
  2. OTP generation and verification processes
  3. SIM lifecycle and network-level metadata
  4. Geolocation and contextual signals
  5. Blockchain-based audit records and policy logs
- Trust boundaries—the system spans multiple trust boundaries, each vulnerable to manipulation:
  1. User device boundary—mobile device, browser, or app interface
  2. Network boundary—mobile operator network, signaling channels, and internet transport
  3. Service boundary—backend servers, AI inference engine, and smart-contract logic
  4. Consensus boundary—blockchain validators participating in identity verification

Ensuring that adversaries cannot cross or exploit these boundaries undetected is central to the design.

- Adversarial capabilities—the threat model considers adversaries capable of the following:
  1. Social engineering telecom operators or service agents
  2. Intercepting or altering network traffic
  3. Spoofing or cloning user devices
  4. Abusing remote eSIM provisioning interfaces to inject fraudulent eSIM profiles or perform unauthorized SIM activations, enabling covert number hijacking without physical SIM access
  5. Manipulating geolocation information or using VPN/location-masking tools
  6. Replaying or injecting fraudulent OTP requests
  7. Exploiting compromised insiders or misconfigured infrastructure
  8. Attempting to tamper with audit logs or off-chain data repositories

  These capabilities inform the specific threat categories considered in the system design.
- Threat categories—based on the assets, boundaries, and attacker capabilities, the system must address the following threat classes:
  1. Identity and credential compromise: Attacks such as SIM swap fraud, phishing, OTP forwarding, and device cloning allow adversaries to impersonate legitimate users or intercept authentication factors.
  2. Spoofing and misrepresentation of context: Adversaries may falsify device identity, geolocation signals, or network attributes to bypass contextual validation.
  3. Tampering with authentication workflows or logs: This includes modifying SIM metadata, altering OTP request flows, or manipulating audit trails stored off-chain or within the blockchain layer.
  4. Information disclosure: Unauthorized access to sensitive metadata, such as mobility patterns, device identifiers, or fraud-detection inputs, may enable profiling or targeted attacks.
  5. Replay and session manipulation: Attackers may reuse prior OTPs, hijack ongoing sessions, or inject forged requests into the authentication pipeline.
  6. Insider misuse or privilege abuse: Malicious or negligent insiders may exploit elevated access to alter policies, modify logs, or leak contextual data.

4.4.2. Security Objectives.

To mitigate the threats identified above, the platform is designed around four primary security objectives:

- Robust user–OTP binding: Ensure that every OTP is cryptographically tied to verified user identity, device integrity, and geolocation context, preventing impersonation and spoofing.
- Early detection of SIM and network anomalies: Detect SIM swap events, number porting, device transitions, or suspicious access behavior prior to OTP issuance.
- Tamper-resistant auditability: Record authentication events, policy decisions, and verification outcomes on an immutable ledger to ensure traceability and support post-incident analysis.
- Privacy-preserving data governance: Minimize exposure of personal data by storing only hashed, anonymized, or reference-derived attributes on-chain, while enforcing user consent and regulatory compliance.

By achieving these objectives, the proposed system can proactively block anomalous OTP requests, detect SIM-swap and location-based fraud in real time, and maintain verifiable, privacy-preserving authentication across multi-operator and cross- border environments.

4.5. Blockchain Layer

The blockchain layer serves as the foundational trust infrastructure for the proposed authentication framework, enabling secure, auditable, and privacy-preserving mobile identity management across multiple MNOs and relying parties. Its core function is to ensure that subscriber identity records and verification outcomes remain tamper-resistant, traceable, and verifiable even when managed by different operators.

The system is implemented on Hyperledger Fabric, chosen for its permissioned architecture, modular consensus model, and suitability for high-throughput telecom environments. A RAFT ordering service provides deterministic consensus and fault tolerance without the computational overhead of proof-of-work mechanisms. Network governance is distributed among participating MNOs, selected relying parties, and an independent identity provider. Each participant operates one or more peer nodes under Membership Service Provider (MSP) control, with digital certificates issued by a consortium-managed Certificate Authority (CA). This decentralized governance structure ensures that no single entity can unilaterally alter or delete identity data.

4.5.1. Data Organization

The blockchain employs a dedicated channel—id-channel—for mobile identity records. This single-channel configuration simplifies coordination, enhances scalability, and ensures a consistent, shared state across authorized peers. The id-channel maintains only minimal, privacy-protected fields:

- msisdn_hash—A salted hash of the subscriber's mobile number. Rotating salts prevent pre-computation or correlation attacks.
- profile_hash—A salted hash referencing the subscriber's encrypted off-chain profile, which may include KYC attributes, device-binding data, and risk metadata.
- operator_id— A unique combination of the Mobile Country Code (MCC) and Mobile Network Code (MNC) identifying the serving operator.

By storing only hashed references and network identifiers, the blockchain avoids holding sensitive PII while maintaining a globally consistent identity index accessible across all participating operators.

4.5.2. Operational Principles

In this architecture, OTP values are never stored on-chain. OTP generation and verification occur at the application layer, with temporary in-memory caching and strict Time-to-Live (TTL) expiry to mitigate replay risks. The blockchain instead records subscriber-state transitions, linking MSISDN and profile hashes for both real-time verification and post-event auditability. Core smart-contract (chaincode) functions include:

- registerUserProfile (msisdn_hash, profile_hash, operator_id)—creates a new or ported subscriber record, supporting seamless number portability.
- updateProfileHash (msisdn_hash, new_profile_hash)—updates profile references when subscriber data changes.
- getLatestProfileHash (msisdn_hash)—retrieves the current profile hash for off-chain identity verification.
- revokeUser (msisdn_hash)/revokeDevice (device_id)—removes identities or devices from active service, enabling fraud remediation and lifecycle control.

4.5.3. Privacy, Compliance, and Security Features

The blockchain layer follows a principle of on-chain minimalism, retaining only the metadata essential for verifiable identity linkage. Key controls include:

- Periodic salt rotation—prevents long-term correlation or dictionary attacks.
- Redactable pointers—leverages chameleon-hash constructions to selectively remove off-chain data references while preserving chain integrity, supporting GDPR-compliant erasure.
- Access control—MSP-based role policies restrict execution and visibility of chaincode operations.
- Hybrid storage model—sensitive attributes (PII, location data, AI features, and KYC files) remain in encrypted off-chain repositories, with the blockchain acting solely as an immutable index and event ledger.

4.5.4. Strategic Value in Multi-MNO Environments

By incorporating the operator_id (MCC/MNC) field, the blockchain layer supports cross-operator identity verification and Mobile-Number Portability (MNP) resolution. When a relying party initiates an authentication request, the framework determines the appropriate serving MNO and retrieves the corresponding profile reference for off-chain validation—without exposing the subscriber's actual MSISDN or other sensitive identifiers. This capability enables scalable, privacy-preserving authentication in multi-operator and roaming environments where interoperability and compliance are essential.

4.6. Geolocation Layer (Contextual Validation)

The geolocation layer links OTP issuance to a validated physical location by correlating multiple independent location signals. The primary source of device location is the serving mobile network operator (MNO), which determines handset coordinates using GSM triangulation or other radio-network positioning techniques (Figure 7). This network-based approach is independent of user-controlled GPS data and reduces susceptibility to spoofing or emulator-based manipulation. The blockchain identity layer identifies the correct MNO through the stored operator identifier (MCC + MNC) associated with the user's MSISDN hash, enabling accurate operator resolution even after number portability.

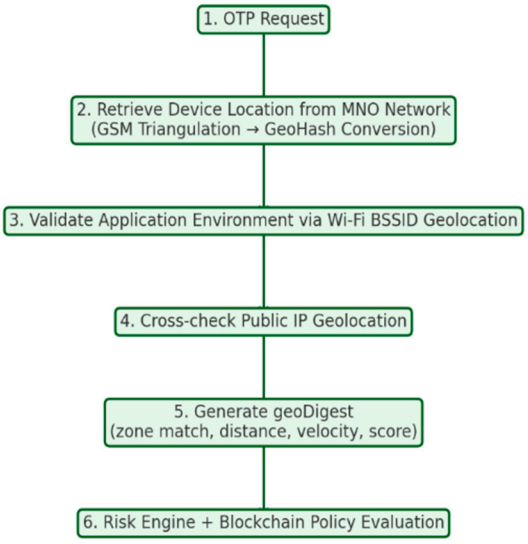

Figure 7. OTP verification flow incorporating cross-validated geolocation signals.

To reduce reliance on a single source and improve robustness, the system also captures Wi-Fi Basic Service Set Identifier (BSSID) information from the environment in which the OTP request originates. When matched against global Wi-Fi geolocation databases, these BSSIDs provide an approximate environmental location. Correlation between Wi-Fi-derived and MNO-derived positions assists in detecting anomalies such as remote access, VPN-based masking, or virtual-device execution.

Public IP geolocation is used as an additional, lower-precision source to support regional consistency checks. Significant discrepancies between IP-based and network-based locations trigger elevated risk (e.g., MNO and Wi-Fi data indicating Paris while the IP address resolves to Singapore).

To maintain privacy, raw geographic coordinates are never stored on-chain. Instead, locations are converted into GeoHash prefixes at precision levels 5–7, which provide an approximate spatial resolution of 4 to 1500 m. These hashed or bucketed representations support location comparison while limiting the risk of user re-identification. As shown in Figures 8 and 9, higher GeoHash precision generally yields higher confidence scores, and levels 5–7 provide a balanced trade-off between privacy and validation accuracy.

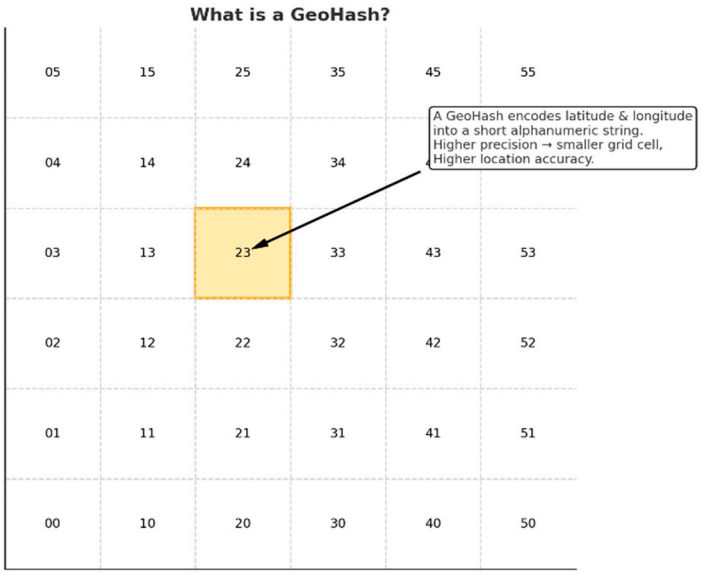

Figure 8. Conceptual representation of GeoHash grid encoding.

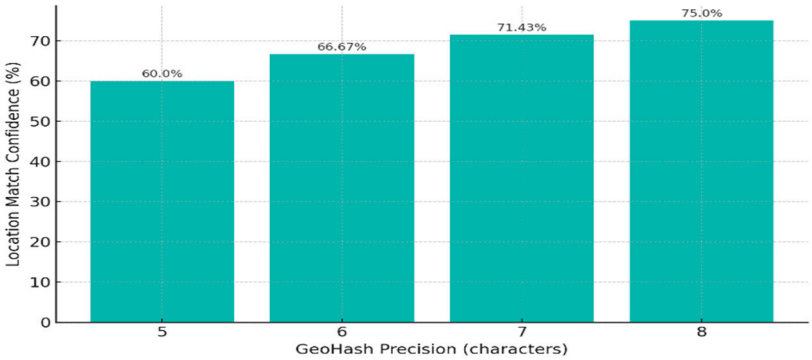

Figure 9. GeoHash location comparison confidence levels.

Although GeoHash levels 5–7 offer suitable granularity for most environments, coarse cell areas in rural regions may still pose identification risks. Future work includes adaptive precision scaling, where smaller cell sizes are applied selectively in densely populated areas to reduce inference risks. Another limitation is the assumption of a trusted prover: the User Identity Management System, which processes raw KYC and location data before ZKP generation. Client-side proof generation and secure-enclave attestations are potential enhancements to minimize the exposure of sensitive data prior to commitment.

The framework additionally assumes the integrity of MNO location data and device GPS modules, but these remain vulnerable to signaling-network exploits (e.g., SS7, LTE) that may allow interception or redirection of location queries. Mitigation strategies include cryptographically signed location attestations using operator-controlled private keys or SIM-based tokens to ensure authenticity. GPS spoofing can also be reduced through multi-source correlation (GPS, cell-ID triangulation, Wi-Fi/Bluetooth beacons) or secure GNSS verification schemes such as TESLA-based broadcast authentication.

After validating signal integrity, geofence policies determine whether the user is within expected operational zones, such as personal trusted locations or service-provider-defined areas (e.g., ATMs, branches, merchant sites). Additional temporal and mobility constraints (e.g., travel velocity, business-hour restrictions, permitted roaming status) refine the evaluation.

The final output of this process is a geoDigest containing zone-match status, distance from the expected location, estimated travel velocity, and a corroboration score reflecting the agreement among MNO, GPS, Wi-Fi, and IP data sources. This digest is forwarded to the blockchain policy engine and the AI-based risk model, enabling OTP decisions that incorporate network-verified and contextually corroborated location information.

4.7. AI Risk Engine (Anomaly and SIM Swap Detection)

The AI Risk Engine forms the analytical backbone of the proposed authentication framework, enabling proactive detection of high-risk events—particularly SIM swap attempts—before the issuance of OTPs. It processes multiple categories of features to generate a comprehensive risk score for each OTP request.

The SIM lifecycle component evaluates indicators such as the time gap between a recent SIM change and the current OTP request, flags for confirmed SIM swaps, and changes to IMSI or ICCID identifiers. Unusually short time intervals, often measured in minutes, may suggest that an attacker has performed a SIM swap and is attempting to capture an OTP before the victim becomes aware.

The device continuity features monitor whether the requesting device remains consistent with historical usage. Changes in International Mobile Equipment Identity (IMEI) or SIM type (e.g., switching from a micro-SIM to a nano-SIM or eSIM) are analyzed in context to differentiate between legitimate upgrades and potential fraudulent behavior.

Behavioral indicators capture OTP request patterns—such as timing, frequency, and destination diversity—along with failed authentication sequences. These patterns help identify abnormal behaviors that deviate from a user's established baseline.

Geo-behavioral analysis compares the location of the SIM change to the OTP request using GeoHash matching, measuring the degree of spatial similarity. A low match score may indicate that the SIM change and OTP request occurred far apart geographically—an anomaly suggestive of fraud. Network intelligence augments this by integrating MNO risk feeds, KYC compliance flags, and known fraud alerts, together with public IP and Autonomous System Number (ASN) classifications. The training features for the AI-based SIM swap detection model are summarized in Table 2, which defines telecom-level and behavioral parameters derived from subscriber activity logs.

Table 2. SIM swap fraud detection model training dataset structure.

Field Name	Description
sim_swap_time_gap_minutes	Time difference (in minutes) between the last SIM change and the current OTP request. A low value (e.g., 1–2 min) may indicate a fraud attempt, as fraudsters typically request OTPs shortly after swapping the SIM.
device_change_flag	Indicates whether the device (e.g., IMEI) has changed compared to historical records. A 'true' value is suspicious if it coincides with a SIM change.
sim_type_change_flag	Indicates whether the SIM type (e.g., physical to eSIM) changed. May signal fraud or a legitimate upgrade—requires correlation with other indicators.
imsi_change_flag	IMSI (Mobile subscriber identifier). A change indicates a SIM swap.
iccid_change_flag	ICCID (SIM card identifier). A change indicates the SIM card was physically replaced.
otp_and_sim_change_geo_hash_length	Shared GeoHash prefix length between SIM change and OTP request locations. A shorter prefix (e.g., 2–4) implies large spatial separation (suspicious); a longer prefix (e.g., 9–10) indicates proximity (normal).
recent_sim_activation_days	Number of days since current SIM activation. Lower values may indicate higher fraud risk.
num_sim_changes_last_30d	Number of SIM changes in the past 30 days. High counts may indicate instability or tampering.
previous_sim_holder_tenure_days	Duration (in days) that the previous SIM was associated with the account. Short tenures may indicate instability or tampering.
account_age_days	Age of the user's mobile account in days. Newer accounts may be more susceptible to fraud.
ip_change_flag	Indicates whether the user's IP address has changed recently.
sim_swap_flag	Label indicating whether the record is classified as a SIM swap fraud attempt (1) or not (0).

The Risk Engine employs a layered modeling approach. Unsupervised anomaly detection methods (e.g., Random Forest-based anomaly scoring) identify novel attack behaviors without prior labeling, while supervised learning techniques are applied when sufficient labeled data exist. Temporal models capture evolving attack patterns over time, allowing for the detection of slow-developing threats. Class imbalance—where fraudulent cases are much rarer than legitimate ones—is mitigated through data weighting strategies and synthetic fraud scenario generation during training.

When an OTP request is made, the Risk Engine returns a real-time risk score via API. This score is then compared against configurable thresholds: low-risk cases proceed automatically, medium-risk cases trigger step-up verification, and high-risk cases are blocked. For transparency and regulatory compliance, key decision factors are logged using Explainable AI (XAI) techniques, ensuring that the reasoning process remains auditable and reviewable.

By combining heterogeneous data sources with a multi-model detection pipeline, the AI Risk Engine achieves high detection accuracy for known threats and strong adaptability to emerging fraud patterns. This approach ensures robust, scalable, and context-aware decision-making for OTP authentication in real time.

4.8. Access and Integration Layer (Microservices)

The Access and Integration Layer orchestrates communication between relying parties (e.g., banks, fintech applications), MNOs, and the underlying identity management components. It provides a secure, service-oriented architecture for mobile identity verification and OTP-based authentication, implemented through RESTful microservices that expose both internal and external APIs for modularity and ease of integration.

4.8.1. Mobile Subscriber Identity Management for MNOs

To facilitate interoperable identity management across telecom providers, the system integrates the Mobile Subscriber Identity Management (MSIM) API, as defined in the mobile-subscriber-identity-management specification. This API allows each MNO to do the following:

- Onboard and manage mobile subscriber profiles.
- Update or retrieve IMSI, ICCID, and MSISDN mappings.
- Record SIM swap or porting events.
- Synchronize metadata such as SIM type and device fingerprint for downstream verification.

Each MNO operates its own blockchain node and communicates with the chain-adapter service to securely hash and store mobile identity profiles. This architecture ensures that subscriber identities remain tamper-evident and auditable while enabling accurate number portability resolution and operator discovery across networks.

4.8.2. Secure Mobile Identity Verification and OTP Authentication

The Secure Mobile Identity Verification and OTP Authentication API, as specified in secure-otp-frameworks, provides a high-assurance interface for real-time subscriber verification and OTP issuance. It integrates seamlessly with external systems such as banking platforms, digital wallets, and online service portals. Key features include the following:

- Pre-OTP verification through KYC validation, SIM change history, and geolocation checks.
- Conditional OTP issuance via decoupled services, ensuring delivery only when risk conditions are met.
- Geo-behavioral validation, including GeoHash matching between SIM events and OTP request locations.
- SIM swap detection based on recent identity changes, device continuity, and telecom metadata.

This API framework enables relying parties to implement adaptive, risk-aware authentication flows, effectively reducing the likelihood of account takeover, SIM hijacking, and session spoofing.

4.8.3. Interconnection Between MNO, Blockchain, and Identity Verification

MNOs interact directly with their private blockchain nodes to write and update hashed user identities tagged with operator-specific metadata (e.g., MCC/MNC). These transactions are routed through the chain-adapter, ensuring that identity anchors are immutable and distributed across all consortium peers. Each profile update—such as a SIM swap—is simultaneously recorded on-chain and made accessible via the MSIM API for real-time verification.

When an OTP request is received from a relying party, the system performs a live identity and location lookup using this module. Based on the MSISDN hash, it identifies the corresponding operator, queries the MNO's identity records, and cross-verifies with

Mobile Number Portability (MNP) mappings and current network context (e.g., cell tower triangulation and SIM issuance metadata).

4.8.4. Northbound and Southbound Integration Support

The platform employs a layered integration model for scalability and interoperability:

- Northbound APIs: Consumed by external services (e.g., banks, fintechs, aggregators). These include the following:
  1. /otp/request—OTP issuance
  2. /otp/verify)—OTP validation
  3. /otp/status/{id}—decision and audit record retrieval
- Southbound APIs: Used by MNOs and internal modules (e.g., geo-service, risk engine). These include:
  1. /mobile-identity/*—subscriber identity queries
  2. SIM change alerts and KYC event notifications
  3. Ingestion endpoints for risk and compliance flags

All APIs are exposed through an auth-gateway that enforces OAuth2 and mTLS authentication, complete with audit logging and observability via Prometheus and OpenTelemetry. End-to-end encryption is maintained, and secrets are protected using Hardware Security Module (HSM)-backed secure vaults to ensure compliance and privacy integrity.

Figure 10 illustrates an example API request for KYC validation during new user registration (e.g., when creating a new bank account). These validations ensure that the registered user's identity is verified before OTP generation.

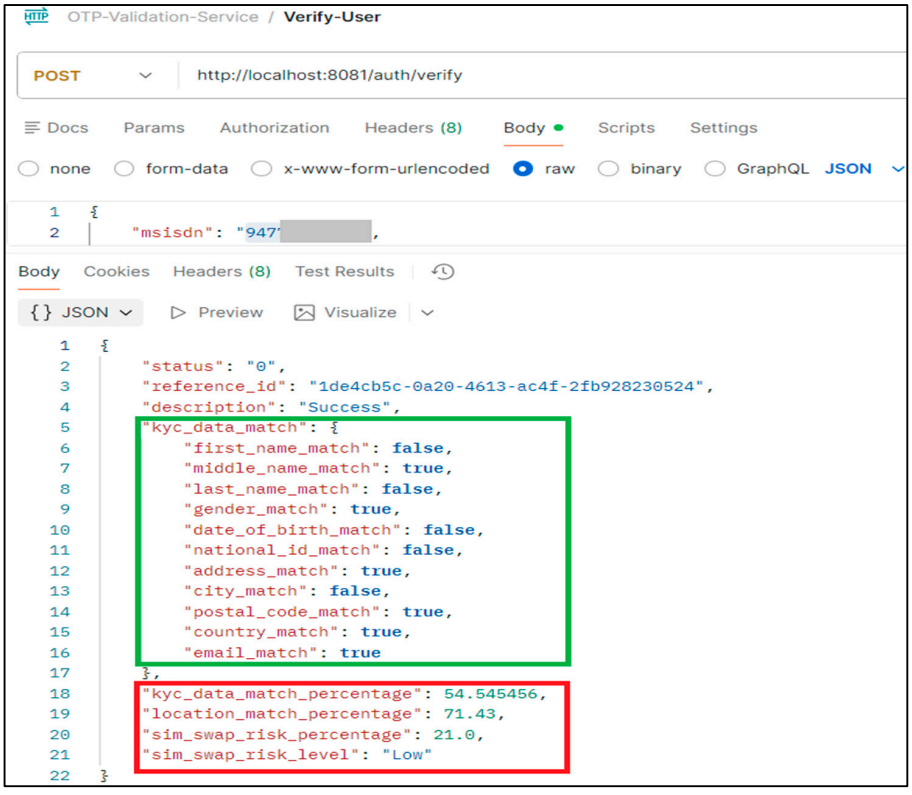

Figure 10. OTP request verification prior to OTP issuance during user registration. The green box shows field-level KYC attribute match results, while the red box presents aggregated percentage scores for KYC similarity, location consistency, and SIM-swap risk.

Figure 11 depicts the OTP request workflow for returning users. In this case, the OTP process omits KYC revalidation but still performs behavioral and contextual checks prior to issuance.

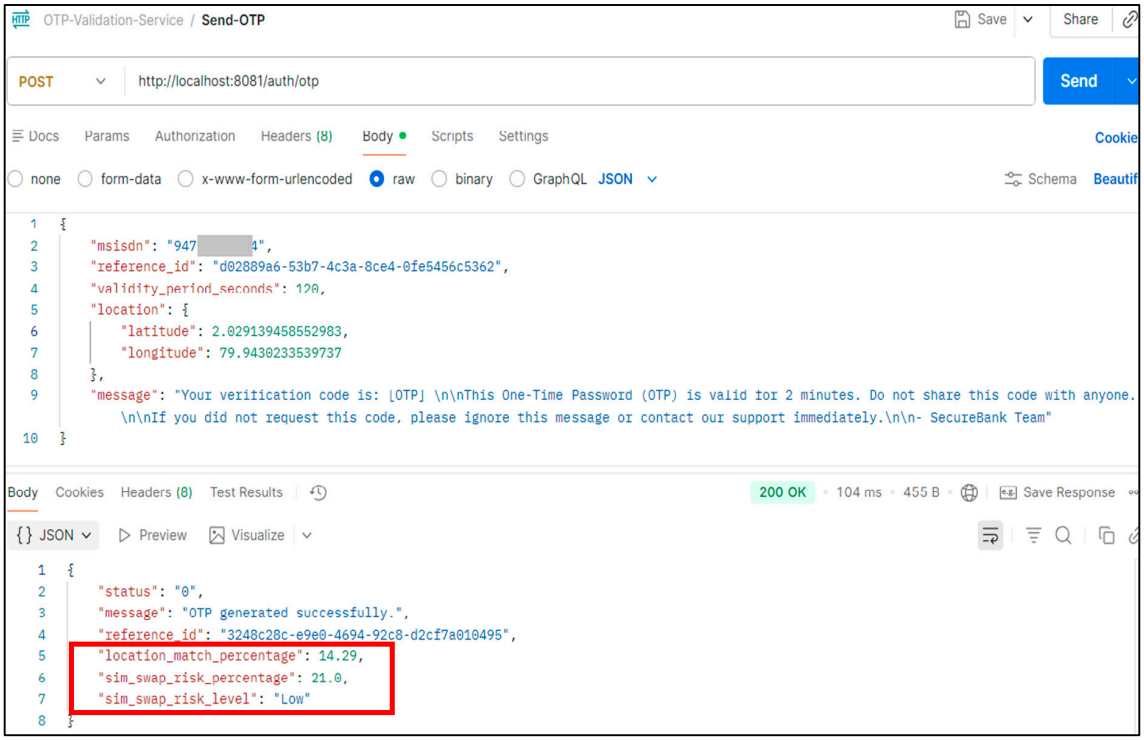

Figure 11. OTP request for verifying a registered user's transaction. The red box shows aggregated evaluation outputs, including the location-match percentage, SIM-swap risk percentage, and the SIM-swap risk level.

All verification processes incorporate geolocation validation. The OTP module interacts with the Mobile Identity Management system to perform real-time KYC and location verification. Figure 12 shows how a user's live location is securely retrieved and compared with the transaction's geolocation, ensuring strong contextual alignment between device location and transaction origin.

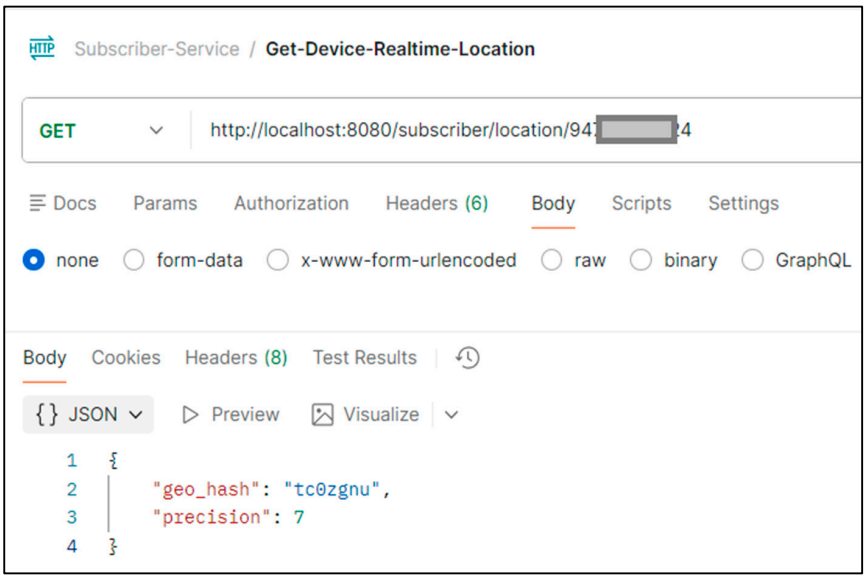

Figure 12. Secure live location retrieval for OTP geo location verification.

4.9. Evaluation Methodology

The evaluation of the proposed secure mobile identity framework was conducted across three primary dimensions: system performance, security effectiveness, and location-based validation accuracy. A combination of blockchain benchmarking, microservice load testing, and AI-based validation techniques was employed to ensure comprehensive analysis.

4.9.1. Performance Benchmarking Tools

Hyperledger Caliper was utilized to assess blockchain network performance under realistic identity-transaction workloads. Metrics such as TPS, block creation and commit latency, and CPU/memory utilization were captured across various configurations. The evaluation compared performance across p50, p95, and p99 percentiles, measuring trade-offs between throughput and latency for diverse endorsement policies. Simulated workloads included mobile subscriber registrations, data write operations, and identity-verification lookups to evaluate end-to-end blockchain efficiency.

Apache JMeter was used to benchmark the API layer, simulating REST-based traffic to measure OTP request and verification performance, registration flow latency, and login session behavior—with and without the AI risk engine enabled. Load scenarios ranged from baseline to peak concurrency, focusing on API throughput, latency, and responsiveness under load.

Additionally, automation testing was performed to validate overall system behavior under simulated end-user interactions. Test cases included sign-up workflows, mobile number registration, and verification across diverse regions and network environments. This ensured consistent responsiveness and decision accuracy throughout the testbed.

4.9.2. AI Model for SIM Swap Fraud Detection

The AI Risk Engine was evaluated using historical SIM lifecycle and network-behavior datasets that included both labeled and unlabeled activity logs. Synthetic anomaly cases, such as simulated SIM-swap events and location spoofing, were added to assess the model's robustness and its ability to adapt to irregular patterns.

Supervised classification methods, including Random Forest, were used to detect behavioral deviations associated with fraudulent activity. Synthetic fraud samples were incorporated into the training and validation sets to benchmark detection capability under controlled conditions. Evaluation metrics included the following:

- Anomaly detection rate, precision, recall, and F1 score.
- Area Under the ROC Curve (AUC) and Precision–Recall AUC (PR-AUC).
- False Negative Rate (FNR), representing undetected SIM-swap incidents.
- False Positive Rate (FPR), representing legitimate cases incorrectly flagged as fraudulent

Model assessment emphasized reducing FNR to limit missed fraud cases while maintaining a low FPR to avoid unnecessary OTP rejections or delays.

4.9.3. Location Accuracy Validation

The geolocation-based validation layer was tested through simulated user behavior in mobile and web contexts. A Location-Based Verification Test was designed to emulate real-world registration, and login attempts from varying distances to assess contextual consistency.

Evaluation involved comparing registered user coordinates against OTP request origins to compute spatial deviations and authentication responses. The AI model's accuracy was measured using confidence scores generated through the GeoHash-based validation

engine, which established a consistent inverse relationship between spatial deviation and confidence level.

OTP requests originating from displaced or spoofed geolocations were automatically rejected when confidence dropped below the adaptive threshold. The geolocation risk engine consistently maintained verification accuracy exceeding 95%, even under high-latency network simulations and varied regional conditions.

4.9.4. Privacy and Compliance Checks

Privacy assurance and regulatory compliance were evaluated alongside system performance. The framework successfully met GDPR-aligned requirements through the following validation tests:

- Data Minimization Audit: Verified that only essential, pseudonymous identity fields were stored or transmitted.
- On-chain vs. Off-chain Boundary Testing: Confirmed sensitive subscriber data remained off-chain, with only hashed references committed to the blockchain.
- Consent Flow Validation: Ensured that UI/UX consent flows properly captured and logged user authorization for data access.
- Retention and Erasure Audits: Validated API and ledger logs for adherence to data-deletion protocols defined under GDPR.

Collectively, these evaluations demonstrate that the proposed authentication framework achieves a balanced integration of performance, fraud resilience, location validation, and privacy assurance—establishing a robust foundation for secure and compliant OTP-based authentication.

5. Results and Discussions

5.1. Evaluation Criteria for Reviewed Studies

To evaluate the reviewed studies in a structured and consistent manner, four criteria were applied. The first criterion identified each study's primary objective, such as enhancing fraud detection, improving identity verification, or strengthening OTP delivery security. It also captured the specific attack vectors addressed, including SIM-swap fraud, phishing, and location spoofing.

The second criterion examined privacy-preserving mechanisms. Studies were assessed based on whether they incorporated techniques such as Secure Multi-Party Computation (SMPC), ZKPs, or encrypted off-chain storage to protect user identity and sensitive data.

Validation methodology formed the third criterion, distinguishing between approaches tested through small-scale simulations, controlled prototypes, or real-world deployments. This differentiation provides insight into the practical maturity of each solution.

The fourth criterion considered integration readiness, which reflects the feasibility of deploying the studied approach in operational environments. Readiness levels were classified as low, medium, or high based on scalability, interoperability, architectural complexity, and implementation maturity.

Together, these criteria support a consistent comparison of prior work, enabling a clearer understanding of how different approaches address OTP security, privacy protection, and deployment feasibility.

5.2. Critical Discussion

The reviewed literature illustrates steady progress in enhancing OTP authentication through the integration of multiple security mechanisms. Recent approaches combine blockchain-based identity management, device-centric verification, and geolocation controls to address threats such as SIM-swap fraud and phishing.

Privacy-preserving techniques, including ZKPs, SMPC, and encrypted or off-chain data storage, align with contemporary data-protection requirements and support secure verification while maintaining auditability. These developments reflect an ongoing shift toward solutions that balance transparency, accountability, and user privacy.

AI-driven anomaly detection contributes additional protection by identifying deviations in behavioral, temporal, or contextual patterns that may indicate fraudulent activity. Adaptive response mechanisms further assist in mitigating identity-based attacks in real time. Together, these techniques advance OTP verification toward more context-aware and predictive threat assessment.

Despite these advances, the scalability, interoperability, and deployment readiness of existing frameworks remain uneven. Many blockchain- and AI-oriented systems have been evaluated primarily through simulations or controlled prototypes due to the limited availability of operational datasets. Comparative findings suggest that performance variations are often driven by environmental factors. These include network latency, model generalization capability, and architectural configuration. Such influences tend to matter more than inherent differences in the underlying methods.

The present research aims to complement, rather than replace, prior approaches by integrating privacy-preserving proofs, a scalable architectural model, and adaptive fraud-detection techniques into a unified framework. Accordingly, the results should be interpreted as context-specific observations within the scope of the experimental environment rather than generalizable performance claims.

Microservice-based architectures and secure API tunneling continue to provide viable pathways for practical deployment within telecom and financial ecosystems. Overall, both the reviewed studies and the proposed framework indicate a gradual evolution toward interoperable, adaptive, and privacy-compliant OTP solutions that can support emerging authentication requirements.

5.3. Experimental Benchmarking Environment

All experiments were conducted on a dedicated workstation equipped with an Intel Core i7-12700 CPU (12 cores/20 threads, 4.9 GHz), 32 GB DDR4 RAM, and an NVIDIA RTX 3060 GPU with 12 GB VRAM, running Ubuntu 20.04 LTS. This configuration provided sufficient capacity for blockchain transaction benchmarking, ZKP proving, and AI model training.

Hyperledger Fabric v2.5 nodes were deployed using Docker containers within a controlled network environment to minimize external latency variation. ZKP proving procedures employed GPU acceleration to support performance evaluation.

Benchmark parameters were selected to reflect representative workloads, including 50,000 OTP authentication and verification requests executed under varying concurrency levels (1, 8, 32, and 128 clients). ZKP verification tests were repeated 50 times per circuit (with 30 k constraints) to measure both p50 and p99 latency distributions.

For fraud-detection benchmarking, the AI model was trained on 100,000 records using an 80:20 train–test split with five-fold cross-validation to reduce overfitting risk. Gradient-boosted trees (100 estimators, depth capped at 5) were evaluated against a logistic regression baseline, with hyperparameters optimized through grid search.

5.4. Dataset Design and Underlying Assumptions

A synthetic dataset comprising 100,000 simulated records was generated using a Python 3.12.0-based framework to support controlled and reproducible evaluation of SIM-swap fraud detection. Synthetic data were required because real telecom datasets involving

SIM-swap incidents are typically inaccessible due to privacy regulations, confidentiality constraints, and the absence of publicly available labeled breach records.

The dataset design followed established practices in fraud analytics and telecommunication risk modeling and was constructed to reflect behavioral and contextual patterns reported in prior empirical studies. This approach enables reproducibility while maintaining alignment with assumptions commonly applied in fraud-detection research.

5.4.1. Dataset Features

Variables were selected to capture characteristics associated with both fraudulent and legitimate OTP interactions. Principal attributes included the following:

- SIM-Swap Time Gap: Interval between SIM activation and OTP request (in minutes) to model elevated fraud likelihood in short activation cycles.
- Identity-Shift Indicators: Changes in device identifiers (IMEI, IMSI, ICCID), reflecting behaviors frequently associated with SIM-swap or account-takeover attempts.
- Geolocation Discrepancy Rate: Degree of mismatch between historical and current user locations to represent potential spoofing.
- Behavioral Features: Number of SIM changes over 30 days, customer age, and account tenure, variables commonly correlated with telecom fraud risk.

A composite fraud score was assigned using weighted indicators—geolocation mismatch (20%), time-based anomaly (30%), device change (25%), and IP-instability (10%)—to classify samples as fraudulent (sim_swap_fraud = 1) or legitimate (sim_swap_fraud = 0). The resulting dataset maintained a balanced class distribution to support stable model evaluation in the absence of public labeled telecom datasets.

5.4.2. Synthetic Anomaly Injection

The AI risk engine was trained and evaluated using datasets augmented with synthetic anomalies derived from anonymized SIM-lifecycle and behavioral patterns. Artificial events such as forced SIM replacements, sudden location shifts, and device-identifier changes were injected to emulate realistic attack traces. Reported performance values (precision = 0.89, recall = 0.88, F1 = 0.88) therefore reflect controlled laboratory testing rather than validation against operational telecom data. These metrics should be interpreted as indicative rather than representative of production environments. Future work aims to collaborate with MNOs to conduct evaluations using anonymized real-world datasets subject to appropriate consent and data-governance controls.

5.4.3. Summary of Associated Benchmark Tables

Table 3 reports end-to-end authentication latency across concurrency levels (1, 8, 32, 128 clients), showing ZKP overhead of approximately +160–230 ms at p50 and +200–270 ms at p95. Median latency remained within 245–380 ms, with p95 values ranging from 310 to 520 ms.

Table 3. End-to-end authentication latency.

Concurrency (Clients)	No-ZKP (p50)	No-ZKP (p95)	ZKP (p50)	ZKP (p95)	Overhead Δp50	Overhead Δp95
1	85 ms	110 ms	245 ms	310 ms	+160 ms	+200 ms
8	95 ms	130 ms	270 ms	350 ms	+175 ms	+220 ms
32	120 ms	180 ms	310 ms	410 ms	+190 ms	+230 ms
128	150 ms	250 ms	380 ms	520 ms	+230 ms	+270 ms

Table 4 summarizes blockchain throughput and commit latency under different endorsement policies. ZKP integration introduced expected computational overhead while maintaining performance levels suitable for real-time authentication.

Table 4. Blockchain transaction performance (Hyperledger Fabric).

Batch Timeout	Max Msg Count	Endorsement Policy	Throughput (TPS)	Commit Latency (p50, ms)	Commit Latency (p95, ms)
2 s	50	OR (MNO1, MNO2)	850 TPS	220 ms	410 ms
1 s	20	AND (MNO1, MNO2)	620 TPS	310 ms	520 ms

Table 5 summarizes the comparative performance of the proposed AI-based fraud detection models—LR and Gradient Boosted Trees (GBT)—against the rule-based baseline system. The metrics highlight substantial improvements in precision, recall, and overall F1-score achieved through supervised learning on telecom-derived features under controlled simulation settings. The AI metrics reported in Table 5 were obtained using the synthetically generated datasets introduced in this section. The results, therefore, represent indicative model performance under controlled experimental conditions rather than validation against real telecom breach data.

Table 5. Fraud detection performance: AI model vs. rule-based baseline.

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC	PR-AUC
Rule-Based Baseline	0.83	0.72	0.61	0.66	—	—
Logistic Regression	0.89	0.84	0.78	0.81	0.91	0.79
Gradient Boosted Trees	0.94	0.89	0.88	0.88	0.96	0.92

Table 6 provides a comparative overview between traditional OTP verification and the proposed ZKP-based OTP authentication. The comparison highlights measurable differences in privacy assurance, auditability, and latency, emphasizing how ZKP integration improves trust and data protection while maintaining acceptable authentication performance.

Table 6. Baseline comparison: traditional OTP vs. ZKP-based OTP authentication.

Aspect	Traditional OTP Verification	ZKP-Based OTP Verification
Data Exposure	Raw KYC attributes and location checks may be visible to the verifier	Only ZKPs revealed; no raw PII exposed
On-Chain Records	None (logs stored centrally)	Audit hashes and proof verification verdicts stored on the blockchain
Authentication Time (median)	85–120 ms	245–310 ms (includes proof generation + chain commit)
Auditability	Limited (central logs may be tampered)	Strong (immutable blockchain ledger)
Privacy Assurance	Moderate	High (sensitive data remains off-chain, hidden in proofs)

Table 7 presents a detailed baseline comparison between the rule-based detection system and the proposed AI-driven fraud detection models (LR and GMT). The results clearly demonstrate the AI models' superior precision, recall, and F1-score, with minimal latency overhead, highlighting their effectiveness in identifying SIM-swap and location-based anomalies under real-time operational constraints.

Table 7. Baseline comparison: AI model vs. rule-based fraud detection.

Aspect	Rule-Based Detection	AI-Based Detection (LR/GBT)
Methodology	Static thresholds on SIM change and location	Supervised learning on multiple telecom features
Precision	0.72	0.84 (LR), 0.89 (GBT)
Recall	0.61	0.78 (LR), 0.88 (GBT)
F1-Score	0.66	0.81 (LR), 0.88 (GBT)
Error Profile	Many false negatives (fraud cases missed)	Stronger recall; significantly fewer missed frauds
Latency (ms)	~0	2–8 (negligible in real-time pipeline)
Overall Assessment	Lightweight but weak coverage of fraud cases	More accurate, robust, and still low-overhead

5.5. Scalability Strategies for Blockchain-Assisted OTP Authentication

A key challenge in blockchain-assisted OTP authentication systems lies in achieving low latency and high throughput without compromising security, privacy, or auditability. To address these challenges, several complementary scalability strategies can be employed:

- Off-chain storage for large or sensitive data: Storing full user profiles, KYC records, or raw geolocation data on-chain would increase storage overhead and verification latency. To mitigate this, only hashed references (e.g., profile_hash, location_hash) are recorded on the blockchain, while sensitive or large data items are maintained in secure off-chain repositories or decentralized file systems. This hybrid model reduces blockchain state growth while retaining verifiable proof anchoring for compliance and auditability.
- Layer-2 blockchain integration: Authentication-related transactions may be delegated to Layer-2 networks (e.g., Polygon, Optimism), where batches of operations are processed off-chain and periodically settled on the main chain. This reduces confirmation time and operational cost while supporting higher throughput, making L2 solutions suitable for real-time authentication workflows.
- ZK-rollups for proof aggregation: ZK-rollup mechanisms allow multiple authentication proofs to be aggregated into a single on-chain transaction. This substantially lowers verification overhead while preserving cryptographic correctness. Such batching improves throughput and cost efficiency, particularly for environments with high request volume.
- Consensus parameter optimization: Consensus settings can be tuned to align with authentication workloads. Adjustments such as shorter block times or optimized endorsement policies (e.g., OR (Org1, Org2) vs. AND (Org1, Org2)) allow the network to handle frequent, low-value verification requests more efficiently. These configurations balance throughput requirements with trust and fault-tolerance considerations.
- Edge-assisted verification: ZKP verification may be delegated to trusted edge nodes located near mobile-network infrastructure. Only the resulting verification records are later committed on-chain. This reduces end-to-end latency and improves scalability, especially in geographically distributed deployments, while maintaining privacy and integrity guarantees.

Together, these strategies offer a modular approach for supporting scalability in blockchain-based OTP authentication systems. They improve operational efficiency and flexibility while maintaining interoperability and compliance across distributed deployments.

5.6. Evaluation Limitations and Future Deployment Considerations

The throughput and latency measurements in this study were obtained using Hyperledger Caliper and Apache JMeter within a controlled, single-region testbed. This environment enabled reproducibility and minimized external network variability but does not reflect the full range of conditions encountered in geo-distributed or multi-operator deployments. In operational settings, factors such as Wide-Area Network latency, propagation delay, bandwidth constraints, and cross-operator consensus synchronization would likely introduce additional delay and variability.

In multi-carrier environments, inter-peer communication overhead and ordering-service coordination may further affect performance. Techniques such as edge-assisted verification and consensus parameter optimization represent potential strategies to mitigate these effects by validating ZKPs closer to mobile-network infrastructure while maintaining a consistent global state. These approaches may support low-latency authentication when deployed across telecom-grade networks.

The experimental findings should therefore be interpreted as indicative performance baselines rather than comprehensive representations of real-world dynamics. Future evaluations, conducted in collaboration with MNOs and using anonymized SIM-swap datasets under appropriate privacy and governance controls, will be required to assess the framework's interoperability, robustness, and fault tolerance under operational conditions.

The AI performance metrics reported in Table 5, derived from synthetic OTP fraud–detection experiments using rule-based, logistic regression, and gradient-boosted tree models, similarly reflect controlled laboratory behavior. Validation using production telecom datasets will be necessary to evaluate model generalization and adaptiveness to real attack patterns.

By presenting results as reproducible benchmarks rather than definitive production outcomes, this study establishes a transparent basis for evaluating blockchain-assisted OTP verification and highlights the need for empirical, collaborative validation to support future deployment in real-world environments.

5.7. Impact of WAN Latency: Geo-Distributed Analytical Emulation

The single-region performance results in Table 4 (850 TPS with a median latency of 220 ms under the OR endorsement policy) provide a baseline under ideal network conditions. However, consortium-based MNO deployments typically operate across multiple geographic regions, where wide-area latency, bandwidth constraints, and ordering delays can affect throughput and commit times. To approximate these conditions, a structured analytical model was developed to assess expected performance in a geo-distributed environment.

The model represents three logical regions (R1, R2, R3), each hosting an MNO peer participating in a Raft-based ordering service. Client transactions originate from R1, consistent with the benchmarking environment described in Section 5.3. Inter-region network behavior is modeled using standard Linux tc netem parameters, applying delay and bandwidth constraints similar to those used in WAN emulation. Three representative latency–bandwidth profiles were examined: 40 ms/200 Mbps, 80 ms/100 Mbps, and 150 ms/50 Mbps.

Hyperledger Fabric parameters were aligned with typical deployment practices, including batch timeouts of 1–2 s, MaxMessageCount values of 20–50, and both OR (Org1, Org2) and AND (Org1, Org2) endorsement policies. Analytical reasoning indicates that commit latency generally increases with inter-region RTT, while throughput decreases as ordering-service and endorsement delays accumulate. The impact is more pronounced

under AND policies, which require multi-peer agreement, whereas OR policies exhibit greater tolerance to latency variation.

This approach constitutes a model-based performance assessment rather than a live distributed emulation. It provides an initial indication of how wide-area latency, endorsement configurations, and batching strategies may influence throughput and transaction-finality times within multi-region authentication networks.

These insights complement the empirical results presented in Sections 5.4–5.6 by extending the analysis beyond controlled single-region conditions. Future work will involve distributed testbed validation across multiple MNO regions to empirically verify the analytical projections and assess system behavior under real operational constraints.

6. Conclusions

This study examined the evolving landscape of OTP security through an analysis of emerging technologies, including blockchain, AI, and geolocation-based authentication. The results highlight the growing importance of strengthening digital identity systems against threats such as SIM swap fraud, phishing, and location spoofing—attack vectors that traditional OTP schemes are not designed to detect or prevent reliably.

The proposed zk-SNARK-based authentication mechanism provides measurable privacy improvements through non-interactive verification; however, it also introduces practical limitations, including the need for a trusted setup, elevated prover-side computational overhead, and constraints on extensibility. These limitations may be addressed in future iterations of the framework. The reported throughput of 850 TPS represents a single-region performance baseline, and future evaluations will extend testing across geo-distributed MNO consortium nodes to account for WAN latency, inter-region consensus delays, and bandwidth constraints. Such evaluations would provide a more representative performance assessment under cross-carrier deployment conditions.

The findings indicate that although decentralized identity management, AI-driven fraud detection, and privacy-preserving geolocation techniques each exhibit potential, most existing implementations operate in isolation. Prior studies often emphasize protocol-specific or detection-specific optimizations but do not present integrated approaches for addressing multi-domain security challenges within a unified framework.

This fragmentation highlights a research gap: the absence of a practically deployable authentication architecture that combines blockchain-based identity verification (including MNP compliance), AI-assisted fraud analytics, and geolocation-based contextual validation. While earlier work has explored privacy-preserving primitives such as ZKPs and geo-hash encoding, performance overheads, regulatory constraints, and limited field testing continue to restrict widespread adoption and operational scalability.

Overall, OTP security in contemporary digital ecosystems remains a multidisciplinary challenge. The combination of privacy-preserving blockchain infrastructures, adaptive AI analysis, and geolocation-based contextual validation represents a potential direction for developing scalable and interoperable authentication mechanisms. Future research may focus on reducing latency, improving interoperability, and advancing the transition of these conceptual models into practical frameworks suitable for deployment within fintech, telecommunications, and digital identity management systems. This study outlines a cross-domain architecture designed to balance performance, decentralization, and privacy in real-world environments.

The proposed ZKP-enhanced OTP authentication framework was evaluated across latency, blockchain throughput, and fraud detection accuracy. The key quantitative observations are summarized below:

- Authentication latency: Traditional OTP verification achieved a median latency of 85–120 ms, while ZKP-enhanced verification required 245–310 ms, resulting in an additional 160–190 ms per request. Total end-to-end latency remained below 0.5 s, which is within typical thresholds for real-time financial and telecom applications.
- Blockchain performance: Using Hyperledger Fabric, the system achieved 850 TPS under an OR (Org1, Org2) endorsement policy with a median commit latency of 220 ms, compared with 620 TPS and 310 ms under the stricter AND (Org1, Org2) policy. This indicates configuration scalability for high-frequency authentication workloads.
- Fraud detection accuracy: The AI model achieved an F1-score of 0.88, with a precision of 0.89 and a recall of 0.88, compared to an F1-score of 0.66 for the rule-based baseline. This demonstrates improved detection of fraudulent activity while maintaining controlled false-positive rates.
- Baseline comparison: Relative to traditional OTP schemes, the proposed system avoids exposing raw KYC or geolocation data by using ZKP-based verification, and the AI component reduces false negatives by more than 40%.

Collectively, these results suggest that the framework provides initial support for enhancing security, privacy, and operational scalability. The ZKP layer adds verifiable privacy protection with manageable latency overhead, and the AI-driven detection component improves resilience to evolving attack patterns. These outcomes provide an initial foundation for potential real-world deployment and for further optimization of decentralized authentication ecosystems.

From a regulatory standpoint, the framework aligns with global data protection requirements such as the GDPR and emerging identity-management standards in telecommunications. By keeping PII off-chain and recording only cryptographic proofs, the architecture supports privacy-by-design principles while enabling auditable accountability. Additionally, the modular integration of blockchain, AI, and geolocation components offers a structure that can be adapted to sector-specific governance requirements, including financial authentication, cross-border interoperability, and user consent management.

7. Future Works

Moving forward, research should focus on developing and testing a unified, scalable model for OTP authentication that integrates private blockchain technology, AI, and geolocation-based validation. The overarching objective is to establish a system that maximizes privacy and security while maintaining low latency. This directly addresses one of the major limitations found in blockchain-assisted authentication systems. This can be achieved through lightweight smart contracts, modular consensus protocols, and enhanced data integrity mechanisms that collectively minimize computational overhead and ensure trust.

Future research will extend performance evaluation to geo-distributed deployments across multiple MNO nodes. This will enable quantitative assessment of network latency, bandwidth constraints, and consensus delays that may occur in real-world consortium environments. Planned experiments include using Hyperledger Fabric's Raft ordering service and edge-assisted ZKP verification to measure end-to-end authentication time and throughput under wide-area network conditions, thereby ensuring scalability and reliability for cross-carrier implementations. Further work will also involve validating the AI risk engine using anonymized real-world telecom datasets to empirically assess detection accuracy, recall, and robustness under operational network conditions. In parallel, upcoming work will focus on mitigating geolocation privacy leakage by applying adaptive

GeoHash precision scaling and exploring client-side or enclave-assisted proof generation to eliminate trust dependencies on centralized provers.

Furthermore, future research will extend the geolocation layer with cryptographic location-attestation protocols and multi-source signal cross-verification to defend against SS7/LTE signaling exploits and GPS spoofing threats observed in real-world telecom environments.

Moreover, future validation will incorporate real-world SIM swap attack traces from publicly accessible datasets such as the FCC Fraud Database, GSMA Intelligence feeds, and anonymized telecom operator logs. These evaluations will benchmark the AI-driven fraud detection model under realistic network and behavioral patterns, allowing calibration of recall–precision trade-offs observed in synthetic simulations. Based on preliminary cross-domain anomaly detection studies, we anticipate a moderate performance drop of approximately 4–7% in F1-score due to increased class imbalance and noise, which will be mitigated through incremental domain adaptation and federated model retraining using privacy-preserving data sharing protocols.

Collectively, these planned extensions will bridge the gap between controlled experimental validation and deployment-scale interoperability, providing a stronger empirical foundation for translating the framework into production-grade telecom environments. Building upon the current framework, future work will implement and evaluate distributed aggregation workflows in Hyperledger Fabric, improving consensus scalability across decentralized nodes. An additional research direction involves developing privacy-preserving data channels and SMPC schemes to maintain encrypted off-chain data protection while supporting federated analytics.

Finally, advancements in AI-driven behavioral intelligence remain a key trajectory. Subsequent research will explore semi-supervised and XAI models for continuous fraud learning, dynamic risk calibration, and model transparency. Integration of temporal geolocation signals, adaptive GeoHash models, and graph-based anomaly detection can further enhance contextual awareness while reducing false alarms. Future investigations should also incorporate cross-sector sandbox evaluations involving telecom operators, fintech institutions, and regulators to validate latency, compliance, and user experience under realistic deployment conditions.

Author Contributions: Conceptualization, G.G.D.S. and D.I.D.S.; methodology, G.G.D.S. and D.I.D.S.; software, G.G.D.S.; validation, G.G.D.S.; writing—original draft preparation, G.G.D.S. and D.I.D.S.; writing—review and editing, G.G.D.S. and D.I.D.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The dataset used to train the SIM swap detection API model was synthetically generated by the authors. It does not contain any personal, sensitive, or real user information. This dataset is available for public access and can be downloaded at: https://github.com/sliit-msc-research/sim-swap-fraud-detection-model-dataset (accessed on 10 October 2025).

Acknowledgments: The authors would like to thank the editor and reviewers for their insightful comments and constructive suggestions, which substantially improved the clarity and quality of this work.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Shynu, T.; Rajest, S.S.; Regin, R.; Steffi, R. OTP as a service in the cloud allows for authentication of multiple services. *Int. J. Orange Technol.* **2023**, *5*, 94–112.

2. Kothinti, K.G. Mitigating one-time passcode (OTP) fraud: Strengthening authentication against emerging threats. *World J. Adv. Res. Rev.* **2025**, *26*, 1368–1378. [CrossRef]
3. Lei, Z.; Nan, Y.; Fratantonio, Y.; Bianchi, A. On the insecurity of SMS one-time password messages against local attackers in modern mobile devices. In Proceedings of the Network and Distributed Systems Security (NDSS) Symposium, Virtual Event, 21–25 February 2021; pp. 1–18. [CrossRef]
4. Goel, A.; Rahulamathavan, Y. A comparative survey of centralised and decentralised identity management systems: Analysing scalability, security, and feasibility. *Futur. Internet* **2024**, *17*, 1. [CrossRef]
5. Khalid, M.I.; Ehsan, I.; Al-Ani, A.K.; Iqbal, J.; Hussain, S.; Ullah, S.S.; Nayab. A comprehensive survey on blockchain-based decentralized storage networks. *IEEE Access* **2023**, *11*, 10995–11015. [CrossRef]
6. Hasan, S.S.U.; Ghani, A.; Daud, A.; Akbar, H.; Khan, M.F. A review on secure authentication mechanisms for mobile security. *Sensors* **2025**, *25*, 700. [CrossRef]
7. Yin, X.; He, J.; Guo, Y.; Han, D.; Li, K.-C.; Castiglione, A. An efficient two-factor authentication scheme based on the Merkle tree. *Sensors* **2020**, *20*, 5735. [CrossRef]
8. Khan, H.U.; Sohail, M.; Nazir, S.; Hussain, T.; Shah, B.; Ali, F. Role of authentication factors in Fin-tech mobile transaction security. *J. Big Data* **2023**, *10*, 138. [CrossRef]
9. Ren, Y.; Leng, Y.; Cheng, Y.; Wang, J. Secure data storage based on blockchain and coding in edge computing. *Math. Biosci. Eng.* **2019**, *16*, 1874–1892. [CrossRef]
10. Kuperberg, M. Blockchain-based identity management: A survey from the enterprise and ecosystem perspective. *IEEE Trans. Eng. Manag.* **2019**, *67*, 1008–1027. [CrossRef]
11. Javaid, N.; Alghamdi, T.A. A comprehensive survey on security, privacy and authentication in blockchain. *Int. J. Web Grid Serv.* **2023**, *19*, 259–286. [CrossRef]
12. McCabe, C.; Mohideen, A.I.C.; Singh, R. A blockchain-based authentication mechanism for enhanced security. *Sensors* **2024**, *24*, 5830. [CrossRef] [PubMed]
13. Lim, S.Y.; Fotsing, P.T.; Almasri, A.; Musa, O.; Kiah, L.M.; Ang, T.F.; Ismail, R. Blockchain technology the identity management and authentication service disruptor: A survey. *Int. J. Adv. Sci. Eng. Inf. Technol.* **2018**, *8*, 1735–1745. [CrossRef]
14. Eren, H.; Karaduman, Ö.; Gençoğlu, M.T. Security challenges and performance trade-offs in on-chain and off-chain blockchain storage: A comprehensive review. *Appl. Sci.* **2025**, *15*, 3225. [CrossRef]
15. Wei, Q.; Li, B.; Chang, W.; Jia, Z.; Shen, Z.; Shao, Z. A survey of blockchain data management systems. *ACM Trans. Embed. Comput. Syst.* **2022**, *21*, 25. [CrossRef]
16. Alghanmi, N.A.; Alghanmi, N.A.; Alghanmi, S.A.; Zhao, M.; Hussain, F.K. Data-driven approach for selection of on-chain vs off-chain carbon credits data storage methods. *Knowledge-Based Syst.* **2024**, *310*, 112871. [CrossRef]
17. Abbas, S.; Abu Talib, M.; Ahmed, A.; Khan, F.; Ahmad, S.; Kim, D.-H. Blockchain-based authentication in Internet of Vehicles: A survey. *Sensors* **2021**, *21*, 7927. [CrossRef]
18. Chandola, V.; Banerjee, A.; Kumar, V. Anomaly detection: A survey. *ACM Comput. Surv.* **2009**, *41*, 15. [CrossRef]
19. Ahmed, M.; Mahmood, A.N.; Hu, J. A survey of network anomaly detection techniques. *J. Netw. Comput. Appl.* **2016**, *60*, 19–31. [CrossRef]
20. Bhuyan, M.H.; Bhattacharyya, D.K.; Kalita, J.K. Network anomaly detection: Methods, systems and tools. *IEEE Commun. Surv. Tutorials* **2013**, *16*, 303–336. [CrossRef]
21. Pourhabibi, T.; Ong, K.-L.; Kam, B.H.; Boo, Y.L. Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decis. Support Syst.* **2020**, *133*, 113303. [CrossRef]
22. Wu, Y.; Dai, H.-N.; Tang, H. Graph neural networks for anomaly detection in industrial Internet of Things. *IEEE Internet Things J.* **2021**, *9*, 9214–9231. [CrossRef]
23. Denœux, T. Behavioral biometrics & continuous user authentication on mobile devices: A survey. *Inf. Fusion* **2021**, *65*, 179–191. [CrossRef]
24. Lim, L.-H.; Ong, L.-Y.; Leow, M.-C. Federated learning for anomaly detection: A systematic review on scalability, adaptability, and benchmarking framework. *Futur. Internet* **2025**, *17*, 375. [CrossRef]
25. Sartayeva, Y.; Chan, H.C.B. A survey on indoor positioning security and privacy. *Comput. Secur.* **2023**, *131*, 103293. [CrossRef]
26. Alawami, M.A.; Kim, H. LocAuth: A fine-grained indoor location-based authentication system using wireless networks characteristics. *Comput. Secur.* **2020**, *89*, 101683. [CrossRef]
27. Oligeri, G.; Sciancalepore, S.; Ibrahim, O.A.; Di Pietro, R. GPS spoofing detection via crowd-sourced information for connected vehicles. *Comput. Netw.* **2022**, *216*, 109230. [CrossRef]
28. Park, S.; Seo, C.; Wang, X.; Lee, Y.; Seo, S.-H. Exclusively in-store: Acoustic location authentication for stationary business devices. *J. Netw. Comput. Appl.* **2024**, *232*, 104028. [CrossRef]
29. Ren, Y.; Li, X.; Miao, Y.; Deng, R.H.; Weng, J.; Ma, S.; Ma, J. DistPreserv: Maintaining user distribution for privacy-preserving location-based services. *IEEE Trans. Mob. Comput.* **2022**, *22*, 3287–3302. [CrossRef]

30. Sahoo, S.S.; Chaurasiya, V.K. Proof of location based delivery system using multi-party virtual state channel: A blockchain model. *J. Supercomput.* **2023**, *80*, 703–733. [CrossRef]
31. Bhutta, M.N.M.; Khwaja, A.A.; Nadeem, A.; Ahmad, H.F.; Khan, M.K.; Hanif, M.A.; Song, H.; Alshamari, M.; Cao, Y. A survey on blockchain technology: Evolution, architecture and security. *IEEE Access* **2021**, *9*, 61048–61073. [CrossRef]
32. Kim, M.; Suh, J.; Kwon, H. A study of the emerging trends in SIM swapping crime and effective countermeasures. In Proceedings of the 2022 IEEE/ACIS 7th International Conference on Big Data, Cloud Computing, and Data Science (BCD), Danang, Vietnam, 4–6 August 2022; pp. 240–245.
33. Kumar, K.; Sihag, V.; Choudhary, G. Geofencing based banking authentication system: A fraud mitigation technique. *Res. Briefs Inf. Commun. Technol. Evol.* **2020**, *6*, 30–40. [CrossRef]
34. Berbecaru, D.G. SAM-PAY: A location-based authentication method for mobile environments. *Electronics* **2025**, *14*, 621. [CrossRef]
35. Roberts, D.; Fisher, J. Securing SIM toolkit-based mobile money applications against SIM swap attacks using user location data. *Int. J. Mobile Secur.* **2023**, *19*, 320–333.
36. Ghaffari, F.; Bertin, E.; Crespi, N. User Profile and Mobile Number Portability for Beyond 5G: Blockchain-based Solution. In Proceedings of the 2023 26th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN), Paris, France, 6–9 March 2023; pp. 187–194.
37. Xu, J.; Xue, K.; Tian, H.; Hong, J.; Wei, D.S.L.; Hong, P. An identity management and authentication scheme based on redactable blockchain for mobile networks. *IEEE Trans. Veh. Technol.* **2020**, *69*, 6688–6698. [CrossRef]
38. Kumar, S. Time-based OTP authentication via secure tunnel (TOAST): A mobile TOTP scheme using TLS seed exchange and encrypted offline keystore. *IEEE Trans. Mobile Comput.* **2023**, *21*, 456–470.
39. Anderson, B.; Lewis, T. A complete one-time passwords (OTP) solution using microservices: A theoretical and practical approach. *Softw. Eng. Rev.* **2023**, *14*, 221–234.
40. Ghaffari, F.; Bertin, E.; Crespi, N. Blockchain-based user profile and mobile number portability for beyond 5G mobile communication networks. In Proceedings of the 2022 4th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS), Paris, France, 27–30 September 2022; pp. 75–78.
41. Nwabuwe, A.; Sanghera, B.; Alade, T.; Olajide, F. Fraud mitigation in attendance monitoring systems using dynamic QR code, geofencing and IMEI technologies. *Int. J. Adv. Comput. Sci. Appl.* **2023**, *14*, 938–945. [CrossRef]
42. Alabdulatif, A.; Samarasinghe, R.; Thilakarathne, N.N. A novel robust geolocation-based multi-factor authentication method for securing ATM payment transactions. *Appl. Sci.* **2023**, *13*, 10743. [CrossRef]
43. Adams, H. GeoMoB: A geolocation-based browser for secured mobile banking. *Mobile Secur. J.* **2023**, *11*, 190–202.
44. Abedi, M.; Yazdifar, M.H.; Parsa, F. Fast location prediction algorithm utilized in enhancing one time password authentication. In Proceedings of the 2012 IEEE Student Conference on Research and Development (SCOReD), Pulau Pinang, Malaysia, 5–6 December 2012; pp. 218–222.
45. Johnson, R. One-time passwords: A literary review of different protocols and their security. *J. Cryptogr. Syst.* **2023**, *9*, 177–189.
46. Komandla, V. Enhancing security and fraud prevention in fintech: Comprehensive strategies for secure online account opening. *Glob. Res. Rev. Bus. Econ.* **2024**, *10*, 102–112. [CrossRef]
47. Cha, B.; Kim, C. Password generation of OTP system using fingerprint features. In Proceedings of the 2008 International Conference on Information Security and Assurance ISA, Busan, Republic of Korea, 24–26 April 2008; pp. 243–247.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.