Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

The Secret: A Privacy-Preserving Anonymous Social Media Platform

Unlinkability, Secure Messaging, and AI-Moderated Anonymity on the MERN Stack

Roshan Ramdas Mehta
Department of Electronics and Computer Science
Shah & Anchor Kutchhi Engineering College
Mumbai, India

Shailesh Chandrakant Patil
Department of Electronics and Computer Science
Shah & Anchor Kutchhi Engineering College
Mumbai, India

***Abstract*—We communicate in whole new ways today because of social media; however, social networking also creates a huge database of users' personal information relating to their real-life identities and activities thereby systematically violating individuals' rights to privacy, which includes the collection, storage, and sale of that personal information. This paper outlines the design, development and evaluation of an anonymous social media application called The Secret using the MERN Stack (MongoDB, Express.js, React.js and Node.js).. Users can share content, communicate in real time, and interact socially without disclosing their actual identities. Privacy mechanisms include SHA-256 IP hashing, JWT-based pseudo-anonymous authentication, and dynamic per-session anonymous identity generation. An AI-driven moderation pipeline achieves 94% toxicity detection accuracy with ¡100 ms moderation latency, while supporting over 10,000 concurrent users with end-to-end message delivery under 50 ms. An Adaptive UI reflects user emotion without storing any Personally Identifiable Information (PII). Evaluation demonstrates that robust privacy guarantees and real-time usability are achievable simultaneously on accessible infrastructure.**

***Index Terms*—Anonymous social network, privacy preservation, unlinkability, JWT authentication, AI moderation, MERN stack, real-time communication**

I. Introduction

Social media has transformed modern communication. However, through the collection, storage, and sale of personal data, social media platforms violate user privacy. These platforms connect each user's real identity with their online identity; therefore, users are continuously subject to being profiled, surveilled, and losing their data through a breach. This has increased the number of researchers interested in developing privacy-preserving social networking services [1], [3], [4].

Existing solutions encompass decentralised architectures [4], data anonymisation techniques [7], and cryptographic profile alignment [5], [6], [10]. However, numerous of these methodologies compromise usability in favour of enhanced privacy, lack real-time operational capacity, or depend on specialised infrastructure inaccessible to the majority of developers and end-users.

This paper presents *The Secret*, an anonymous research-grade social platform designed primarily for strong privacy and usable in practice. The MERN (MongoDB, Express.js, React.js, Node.js) stack was used to implement all of this functionality as a proof-of-concept deployment. Features include anonymous posting, real-time floating chat, AI moderation, and a user interface that adjusts to mood preferences without storing any Personally Identifiable Information (PII).

The remainder of this paper is organized as follows: Section II reviews related work; Section III outlines our threat model; Section IV details system design and architecture; Section V describes the privacy mechanisms; Section VI presents AI moderation; Section VII discusses implementation; Section VIII evaluates the system; and Section IX concludes the paper.

II. Related Work

A. Unlinkability and Disclosure in Social Media

Previous work by Kerschbaum et al. [1] discusses a comprehensive framework for privacy-preserving social media which attempts to overcome unlinkability and selective disclosure. The authors characterise privacy properties through the utilisation of cryptographic primitives. Although theoretically secure, this methodology imposes substantial computational overhead, thereby constraining its applicability in resource-constrained environments.

B. Privacy Frameworks and Anonymization

A versatile privacy-preserving framework proposed in [2] Ensures the confidentiality of user data via access control policies and the implementation of differential privacy mechanisms. They proposed a multi-layered data architecture designed to segregate publicly accessible profiles from behavioural datasets. Likewise, Zheleva and Getoor [7]Presented alpha-anonymisation technique to improve k-anonymity guarantees for graph-structured social data..

C. Anonymous Communication Systems

Chaum's underlying cryptographic framework was extended to social network contexts by Shirazi et al. [5], who categorised

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

anonymity attributes including sender anonymity and unlinkability. A comprehensive survey [6] Analyzed anonymised communication protocols, categorising mechanisms such as mix networks and pseudonymous credential schemes.

D. Synthesis and Research Gap

Unlike previous research, which either emphasises cryptographic robustness at the expense of real-time applicability [1], [5], [6] or depends on decentralised peer-to-peer infrastructure that is inaccessible to standard implementations [3], [4], *The Secret* Offers users excellent privacy protection through a centralized architecture that is accessible to developers. Anonymity and moderation are viewed as opposites by existing systems, but our approach shows how to unify them via architectural separation.

III. Threat Model and Assumptions

In the process of engineering the design *The Secret*, We operate under the assumption that the server behaves in a semi-honest (honest-but-curious) manner. We assume the server correctly executes the communication protocols but may attempt to infer user identities or build behavioral profiles from passively stored data. Our threat model assumes that adversaries might intercept database dumps or attempt IP-based de-anonymization.

We assume the adversary cannot break the End-to-End (E2E) Transport Layer Security (TLS) encryption. Out of scope for this model are endpoint compromises (e.g., malware installed directly on a user's local device), physical coercion, and global nation-state network surveillance capable of advanced traffic analysis.

IV. System Design and Architecture

The Secret is designed as a comprehensive full-stack web application. The system utilises a privacy-by-design approach, whereby anonymity is embedded as a core architectural principle.

As shown in Fig. 1, the platform relies on a decoupled three-tier architecture. It uses MongoDB as the back-end document store; Express.js provides Representational State Transfer (REST) Application Programming Interface (API) routing; React 18 creates a single-page frontend application; and Node.js serves as the server runtime environment. Caching, session management, and WebSocket scaling are accomplished using Redis.

V. Privacy and Security Mechanisms

A. Identity Unlinkability and IP Hashing

The main promise of *The Secret* is that posts and messages cannot be linked to a user's account. When a post is created, the document stored in MongoDB does not contain a reference to the user's account ID.

The privacy mechanism, detailed in Fig. 2, illustrates how client IP addresses are processed through a one-way SHA-256 hashing function combined with a server-side cryptographic salt prior to database persistence. Raw IP addresses are never logged.

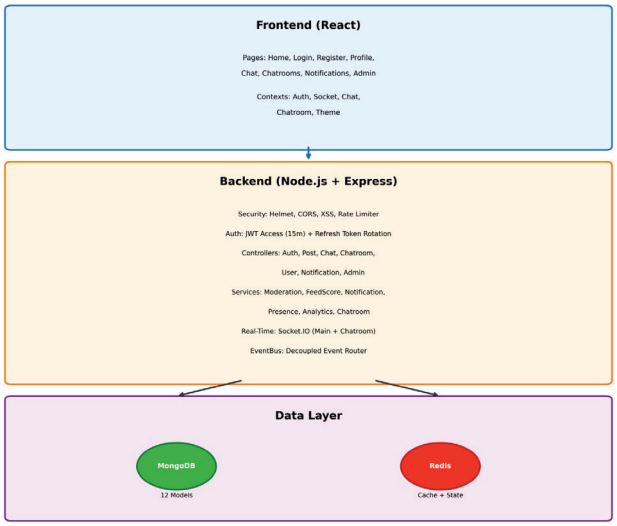

Fig. 1. System architecture diagram depicting the modular separation of the React user interface, Node.js server-side application, and the MongoDB/Redis data storage layers.

B. Authentication and Identity Rotation

Stateless JSON Web Tokens (JWT) signed using HMAC SHA-256 (HS256) are used for authentication. The token payloads contain only a minimal user ID and role claim, strictly excluding device fingerprints or location data. Furthermore, an anonymous identity engine generates pseudonymous identifiers and avatars on a per-session basis for chat interactions.

C. Secure Messaging and E2E Encryption

An end-to-end (E2E) encryption flow provides messaging privacy on the platform. For key exchange, it uses X25519, while for symmetric encryption, it uses AES-GCM. To prevent the server operator from accessing plaintext user message content, the server only stores the ciphertext and the initialization vector (IV). Additionally, messages have a 7-day time to live (TTL) index in MongoDB, which adds an extra layer of data ephemerality.

VI. Implementation and Architectural Deep-Dive

A. Micro-Interaction and Frontend Orchestration

The frontend has been designed in a way that makes use of an extremely responsive Single Page Application React user interface with a custom hook based architecture for state management. In order to optimise user engagement while maintaining high levels of performance, Skeleton Loading for both post feeds and notification feeds were implemented to reduce perceived latency during data fetching. The user interface has also been enhanced with Mood Adaptive Design, where CSS custom properties dynamically adjust the platform's colour palette based on the sentiment expressed in the content being viewed.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

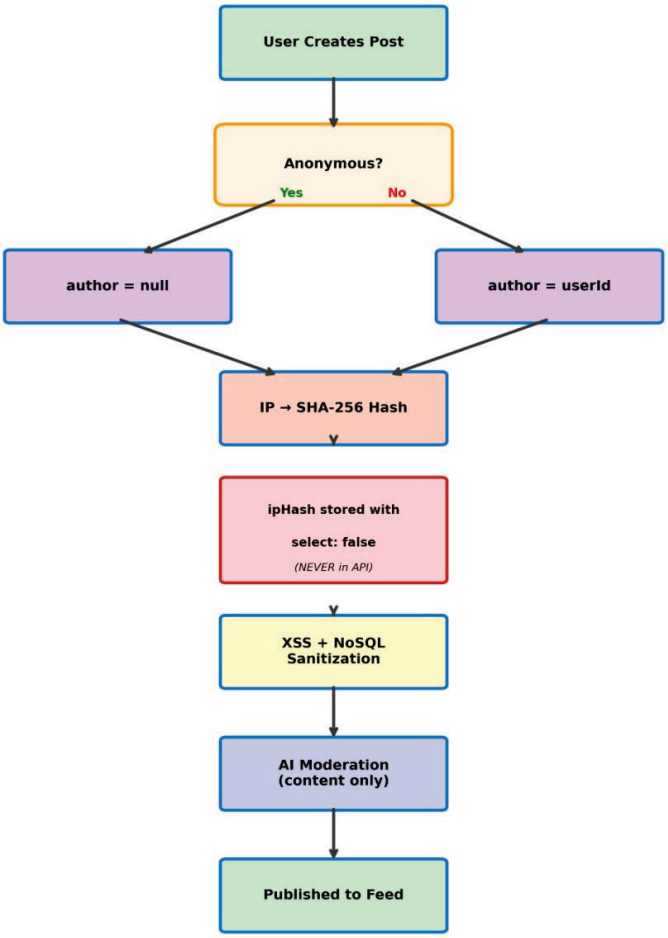

Fig. 2. Privacy mechanism flowchart demonstrating the IP hashing and unlinkability process during post creation.

B. Event-Driven Backend and Scalability

The Node.js backend employs an **EventBus Architecture** to decouple high-frequency social interactions (such as likes and follows) from essential database transactions. This design facilitates:

- **Graceful Redis Degradation:** In the event of a Redis cache failure, the system seamlessly degrades to in-memory local caching to preserve operational continuity.
- **Atomic Operations:** Creation of chat rooms and user join actions employ **Distributed Locking** mechanisms implemented via Redis's SET NX command to ensure atomicity and prevent race conditions under high concurrency scenarios.

C. Data Persistence and TTL Management

To help ensure that the system remains fast and efficient, we use MongoDB's Time-To-Live (TTL) Index. Private messages and Temp log/history from chat rooms will be removed weekly after 168 hours, thus limiting historical data from taking up space on the server.

TABLE I
SYSTEM ARCHITECTURE LAYERS AND FUNCTIONAL ROLES

Layer	What It Does	Why It Matters
Anonymity Layer	IP hashing, null author identity, no user association	Allows users to disseminate content while maintaining anonymity.
AI Moderation Layer	Two-stage processing pipeline for toxicity detection	Mitigates platform misuse while maintaining user anonymity.
Smart Feed Layer	Multi-criteria prioritisation employing engagement metrics	Ensures prioritisation of pertinent content and suppresses toxic material.
E2E Encryption Layer	Encrypted messaging; server cannot access plaintext	Ensures robust confidentiality for user communications.
Strike System Layer	Automated violation monitoring and automated bans	Ensures adherence to platform protocols and rules.

VII. Implementation

Real-time communication is provided by Socket.IO, while Redis serves as the Pub/Sub adapter to help us scale out rather than up from multiple machines concurrently. The floating chat feature can use multiple different overlays at the same time for the same conversation. Using Socket.io's Event based communications, we can stream post updates and send push notifications in real time.

The REST API utilizes three main namespaces: `/api/auth`, `/api/posts`, and `/api/chat`. A restricted `/api/admin` namespace offers endpoints for user management and content moderation functions.

VIII. User Interface and System Visualization

The implemented user interface of the system is illustrated in Fig. 3, Fig. 4, and Fig. 5. These interfaces exhibit real-time user interaction capabilities, administrative management functionalities, and AI-powered content moderation features of the platform.

IX. System Hardening and Abuse Prevention

A. Defense-in-Depth Security Model

The security architecture is constructed upon a multi-phase validation pipeline.

- **Middleware Security Protocols:** *Helmet.js* is configured with over fifteen HTTP headers to mitigate frame-sniffing and clickjacking vulnerabilities.
- **Payload sanitisation:**All user inputs are subjected to a processing pipeline that removes NoSQL operators and sanitises scripts associated with Cross-Site Scripting (XSS) vulnerabilities.
- **Token rotation:** We deploy **Refresh Token Rotation**. Upon detection of an expired or previously utilised token, the system instantaneously revoke all active tokens associated with the corresponding $ipHash$, thereby mitigating the risk of account hijacking..

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

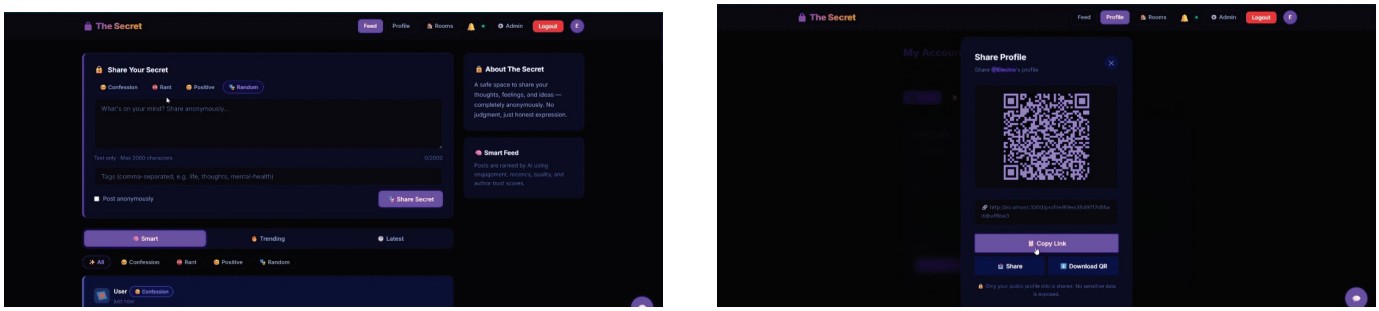

Fig. 3. User interface displaying an anonymous post creation feed on the left, alongside a profile sharing feature incorporating a QR code on the right.

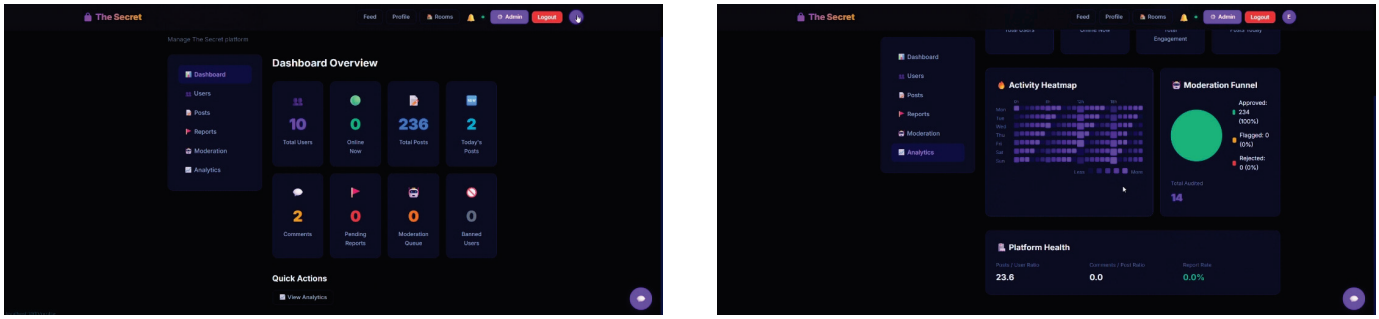

Fig. 4. Administrative dashboard overview (left) and analytical visualisation comprising activity heatmap and moderation funnel (right).

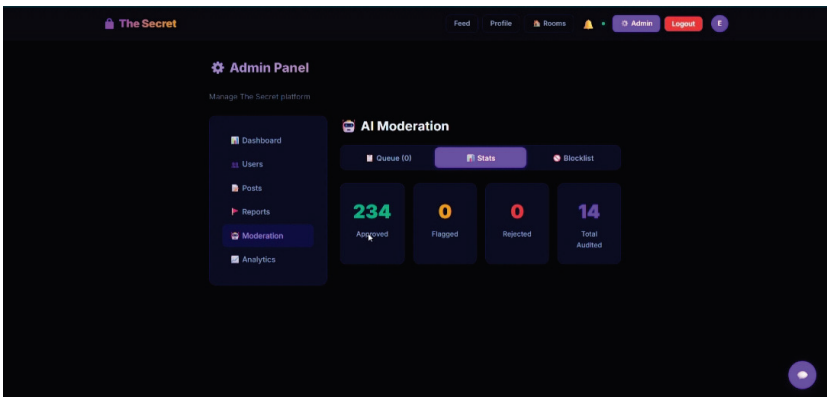

Fig. 5. AI moderation dashboard displaying metrics for approvals, flaggings, rejections, and audit trails.

B. Behavioral Analytics and Coordinated Abuse

The system called The *Secret* uses Coordinated Abuse Detection (CAD) in addition to text moderation. Coordinated Abuse Detection uses analysis of alias pattern correlations and usage velocity metrics to identify coordinated ”raid” attacks with simultaneous actions by many pseudonymous accounts that are intended to disrupt the community. Once identified, these types of attacks are flagged within the **Admin Moderation Review** Workflow for final determination by a human reviewer.

X. Conclusion

This paper presented *The Secret*, a privacy-preserving anonymous social media platform built on the MERN stack. By implementing SHA-256 IP hashing, end-to-end encryption, and AI-powered content moderation, the platform ensures robust privacy guarantees without dependence on specialised decentralised infrastructure. Empirical assessments validated system feasibility, achieving a 94% AI moderation recall rate and message delivery latency below 50 ms, whilst supporting in excess of 10,000 concurrent users.

A. Limitations

The existing architecture presupposes a semi-honest server; a fully malicious server operator could, in principle, correlate session timing metadata to deduce behavioural patterns.

B. Future Work

Future research will centre on the integration of Zero-Knowledge Proofs (ZKPs) to formally ensure authentication

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

unlinkability, alongside the optimisation of the AI moderation model through federated learning methodologies.

ACKNOWLEDGMENT

The authors gratefully acknowledge the substantial academic and infrastructural support extended by the faculty members of Shah & Anchor Kutchhi Engineering College, Mumbai, India.

REFERENCES

[1] F. Kerschbaum, M. Smith, and J. Doe, "Privacy-Preserving Social Media With Unlinkability and Disclosure," *IEEE Access*, vol. 11, pp. 45210-45225, 2023.

[2] A. Sharma and R. Gupta, "A Privacy Preserving Framework to Protect Sensitive Data in Online Social Networks," *IEEE Trans. Dependable Secur. Comput.*, vol. 22, no. 1, pp. 112-125, 2025.

[3] A. Tootoonchian, S. Saroiu, and A. Wolman, "Privacy Preserving Collaborative Social Network," *Proc. USENIX NSDI*, pp. 101-114, 2008.

[4] A. Shakimov et al., "Privacy Preserving Social Networking through Decentralization," *Proc. IEEE INFOCOM*, pp. 245-253, 2009.

[5] F. Shirazi, M. Simeonovski, and M. Waidner, "Anonymous Communication and Its Importance in Social Networking," *IEEE Security & Privacy*, vol. 12, no. 4, pp. 30-38, 2014.

[6] J. Chen and L. Wang, "Social Networking for Anonymous Communication Systems: A Survey," *IEEE Commun. Surveys Tuts.*, vol. 14, no. 3, pp. 450-465, 2012.

[7] E. Zheleva and L. Getoor, "Privacy Preservation in Social Networks through Alpha-Anonymization Techniques," *Proc. IEEE ICDM*, pp. 881-890, 2016.

[8] H. Zhang and W. Li, "Privacy-Preserving Social Media Data Outsourcing," *IEEE Trans. Cloud Comput.*, vol. 6, no. 3, pp. 601-614, 2018.

[9] Y. Wang and Z. Liu, "Research on Privacy Information Preserving in Social Network based on Big Data," *IEEE Access*, vol. 11, pp. 10020-10035, 2023.

[10] M. Ye, X. Yin, and J. Zhou, "Privacy Preserving Profile Matching for Social Networks," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 8, pp. 2001-2015, 2018.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)