

An Adaptive Continuous Authentication Framework Using Zero-Knowledge Proofs for Session Hijacking Prevention

Jayabandi Jersan¹, Prof. Dhavala Lalitha Bhaskari²

¹Department of Computer Science and System Engineering, Andhra University, Visakhapatnam, India

²Department of Computer Science and System Engineering, Andhra University, Visakhapatnam, India

Abstract—Conventional web authentication typically establishes trust during an initial login challenge and subsequently relies on a persistent session cookie or bearer authorization token until explicit expiry or manual logout. In decentralized Web3 ecosystems and cloud platforms, a compromised or stolen session credential may therefore be exploited repeatedly without requiring proof that the legitimate identity holder remains active at the endpoint.

The authors' previously published implementation combined client-side recursive zero-knowledge proofs (zk-SNARKs), freshness nonces, browser-native WebAssembly (Wasm), and read-only Layer-2 verification using a fixed 60-second heartbeat. This paper substantially extends that foundational paradigm by designing an adaptive, risk-aware continuous authentication architecture: routine cryptographic checks remain unobtrusive and computationally lightweight for stable sessions, whereas anomalous behavioral signals or network discrepancies dynamically shorten the re-verification interval, and critical security breaches prompt instantaneous verification or immediate session revocation.

Throughout this adaptive verification cycle, private witness credentials reside strictly within client volatile memory closures, guaranteeing zero exposure of raw credentials across network boundaries while preserving algebraic circuit soundness. This manuscript articulates the end-to-end framework architecture, formal risk-state policy formulations, threat vector mitigation models, an analytical comparison of scheduled computational proof counts, and a reproducible empirical evaluation methodology.

The resulting model bridges the historic trade-off between uncompromising Zero-Trust security and frictionless client performance, providing a scalable, gas-free identity layer for modern decentralized Web3 applications.

Index Terms—Adaptive authentication, continuous authentication, zero-knowledge proofs, session hijacking, risk-based authentication, zero-trust architecture, Groth16 zk-SNARKs, WebAssembly, Polygon Amoy, Web3 security.

I. INTRODUCTION

In contemporary distributed computer systems, web applications and decentralized application (dApp) frameworks commonly issue a persistent session cookie or cryptographic bearer token (such as a JSON Web Token) upon successful completion of an initial login challenge. While this perimeter defense paradigm ensures initial operational convenience, the resulting security posture presumes that the operating client endpoint remains permanently benign for the duration of the token's lifetime. However, empirical threat analyses demonstrate that the vast majority of identity compromises occur post-authentication: adversaries exfiltrate active session tokens through cross-site scripting (XSS), browser infostealers, man-in-the-middle network interception, or unauthorized physical workstation access when an operator temporarily steps away [1], [2].

Under conventional Identity and Access Management (IAM) systems, the server infrastructure remains oblivious to post-login session compromise until token expiration or manual user revocation [3]. To mitigate these systemic vulnerabilities, modern security doctrine mandates Zero-Trust Architecture (ZTA), strictly adhering to the axiom 'Never Trust, Always Verify' [2]. Enforcing continuous identity verification directly within decentralized Web3 environments, however, introduces acute economic and computational bottlenecks. Broadcasting recurring

cryptographic identity checks as state-changing transactions over public blockchain networks results in prohibitive gas expenditures, cumulative transaction settlement lag, and repetitive wallet signature pop-ups that degrade the user experience [4], [6].

The authors' earlier research published in QUEST-AI 2026 introduced a continuous-authentication implementation utilizing recursive Zero-Knowledge Proofs (zk-SNARKs), dynamic epoch freshness nonces, client-side WebAssembly execution, and asynchronous static contract reads (eth_call) on the Polygon Amoy Layer-2 testnet [1]. That baseline architecture achieved sub-second client proving latency and verified session continuity with exactly zero gas expenditure (0 wei). Nevertheless, that foundational system operated across a rigid, uniform 60-second heartbeat schedule regardless of runtime telemetry or environmental anomalies.

This paper extends that research by proposing an adaptive scheduling and risk-assessment architecture. Instead of enforcing an inflexible verification interval across all session lifecycles, our adaptive framework monitors security-relevant client telemetry (such as tab refocus events, Web3 wallet connectivity shifts, IP address discrepancies, and failed proof events) and dynamically adjusts the proof frequency between low, medium, high, and critical states. Furthermore, critical session breaches bypass scheduled delays entirely to trigger instantaneous on-chain verification or local token invalidation in sub-100 ms timeframes, thereby optimizing client battery and CPU load while drastically compressing the exposure window for hijacked sessions.

The major technical contributions of this research are as follows: (1) An adaptive policy orchestration layer that dynamically modulates proof generation intervals according to ambient session risk; (2) Preservation of strict client-side witness privacy where private credentials never leave volatile WebAssembly memory; (3) A zero-gas Layer-2 verification engine utilizing off-chain static view calls to execute bilinear pairing checks on Polygon Amoy; (4) An analytical modeling framework evaluating computational proof count and data transmission over multi-tier operating scenarios; and (5) A rigorous, reproducible experimental methodology establishing baseline benchmarks across workstation, laptop, and mobile device architectures.

II. RELATED WORK AND THEORETICAL FOUNDATION

A. Zero-Trust and Continuous Authentication Paradigms

Traditional perimeter-based access control models establish an implicit trust boundary around authenticated network clients [2]. In contrast, Zero-Trust Architecture treats every internal and external request as potentially hostile, requiring continuous evaluation of device health, credential legitimacy, and user presence [6], [8]. Early continuous authentication frameworks attempted to verify presence via behavioral biometrics, including keystroke dynamics, mouse trajectory tracking, and continuous webcam facial recognition [8], [9]. Although behavioral biometrics reduce manual user interaction, streaming fine-grained sensor telemetry to remote cloud servers introduces severe privacy infringements, data protection liabilities, and vulnerability to telemetry replay attacks [5], [10].

In decentralized Web3 networks, continuous verification is further complicated by consensus mechanics. Conventional blockchain-based Self-Sovereign Identity (SSI) frameworks require users to cryptographically sign challenge strings for each resource access [11]-[13]. Submitting these signed challenges to on-chain smart contracts incurs cumulative transaction fees (gas) and introduces network confirmation latencies that render millisecond-scale session management computationally and economically infeasible [7], [14].

B. Zero-Knowledge Proofs (zk-SNARKs) and Privacy Preservation

Zero-Knowledge Proofs (ZKPs) allow a prover to mathematically demonstrate possession of a secret witness satisfying an algebraic constraint system without exposing the secret itself [1], [15]. Non-interactive arguments of knowledge (zk-SNARKs), particularly the Groth16 protocol, generate succinct, constant-size proofs (comprising three elliptic curve group elements) that can be verified on-chain in milliseconds [1]. In decentralized identity architectures, ZKPs decouple authentication from credential disclosure: a user can prove possession of a private key matching an on-chain public identity commitment without broadcasting their private key across public RPC networks [3], [16].

A central technical barrier to deploying zk-SNARKs in browser runtimes is circuit constraint complexity. Standard hash algorithms such as SHA-256 and Keccak-256 are optimized for bitwise arithmetic (AND, XOR, rotations), resulting in 25,000 to 148,000 Rank-1 Constraint System (R1CS) constraints per hash computation. In contrast, the Poseidon hash function is specifically engineered over prime fields (such as the BN254 scalar field), utilizing algebraic S-boxes (x^5) that compress cryptographic commitments down to exactly 240 R1CS constraints [1], [10]. This mathematical optimization enables consumer web browsers to synthesize valid cryptographic proofs in sub-second intervals.

C. Layer-2 Verification and Read-Only State Queries
Executing zk-SNARK verification on public blockchain ledgers typically involves deploying an EVM-compatible Verifier smart contract that evaluates bilinear pairing equations [1], [5]. Under standard transaction pipelines (`eth_sendRawTransaction`), miners process verification calls as state transitions, consuming computational gas and requiring native network token fees [6]. However, in session authentication contexts where the goal is verification rather than permanent ledger modification, verification calls can be dispatched as asynchronous static calls (`eth_call`) [1].

Static view calls instruct an RPC node to execute the Verifier contract's pairing logic off-chain within local node EVM runtime, returning a boolean verification outcome (true/false) in real time without broadcasting a ledger transaction. This architecture eliminates transactional gas tolls (yielding an effective cost of 0 wei / \$0.00 USD) while inheriting the decentralized security and immutable verifying key anchored on the Layer-2 network [1], [7].

III. PROBLEM STATEMENT AND OBJECTIVES

A. Limitations of Static and Uniform Verification

While the authors' prior implementation demonstrated the technical viability of zero-gas continuous verification via fixed 60-second heartbeats [1],

deploying a uniform verification schedule across diverse operating conditions exhibits significant structural limitations: (1) **Resource Inefficiency:** In tranquil, low-risk operating scenarios, synthesizing cryptographic proofs every 60 seconds generates redundant CPU cycles, increases battery depletion on mobile hardware, and floods RPC infrastructure with unneeded requests; (2) **Temporal Vulnerability Gap:** Conversely, during elevated-risk conditions, a 60-second fixed interval leaves an unacceptably wide window of exposure during which an adversary possessing an exfiltrated token can execute unauthorized operations before the next scheduled check [1], [4].

B. Research Objectives

To overcome these limitations, this research formulates an adaptive, risk-responsive continuous authentication framework. The concrete objectives are: (1) To formalize a multi-tier session-risk policy model defining Low, Medium, High, and Critical security states; (2) To engineer an autonomous client-side scheduler that maps dynamic session risk to variable proof generation frequencies (60s, 30s, 10s, and immediate); (3) To maintain strict client-side witness privacy where private credentials never leave volatile WebAssembly memory; (4) To establish immediate out-of-band verification and sub-100 ms revocation triggers for critical fault events; and (5) To construct a reproducible analytical and experimental validation framework comparing fixed and adaptive verification performance.

IV. PROPOSED SYSTEM ARCHITECTURE

The proposed framework decouples the cryptographic proving mechanism from the behavioral policy engine. The overall architecture spans six cohesive functional layers: (1) Session Initialization, (2) Client-Side Event Monitoring, (3) Risk Evaluation Engine, (4) Adaptive Scheduling Loop, (5) Browser-Native Proving Engine, and (6) On-Chain Layer-2 Settlement. Below, Fig. 1 illustrates the complete system architecture, operational components, and data interaction.




Fig. 1. End-to-end working architecture of the adaptive continuous authentication framework: client event telemetry, adaptive risk-scoring scheduler, WebAssembly Groth16 prover, and Polygon Amoy zero-gas Layer-2 verification engine.

A. Detailed Component Interaction

As illustrated in Fig. 1, the client browser environment establishes an isolated execution sandbox utilizing WebAssembly (Wasm). When the user completes initial wallet binding via MetaMask, the private secret witness s is loaded strictly into volatile Wasm heap memory closures, bypassing unencrypted browser persistent storage (such as localStorage, sessionStorage, or IndexedDB) to eliminate cross-site scripting (XSS) exfiltration risks [1]. Simultaneously, the master public commitment $H_{pub} = \text{Poseidon}(s)$ is anchored within the deployed Verifier contract on Polygon Amoy.

During active operation, the Risk Monitor continuously observes ambient session telemetry and feeds normalized event indicators into the Adaptive Scheduler. Based on the calculated risk score, the scheduler determines the next verification deadline. When the heartbeat timer elapses, the client snarkjs prover compiles the private secret s , dynamic epoch nonce n_t , and previous state into a Groth16 proof vector $\text{argv} = [\pi_A, \pi_B, \pi_C, \gamma_{pub}]$ in sub-second latency. This payload is transmitted via ethers.js to the Polygon Amoy Verifier contract (0x548d...5E6) through an asynchronous eth_call static view. If the

pairing checks balance, the session status remains 'Verified: Access Granted'. If validation fails or a critical disconnect occurs, session tokens are expunged in under 100 ms.

B. Mathematical Formulation of Adaptive Risk Assessment

The risk evaluation engine calculates a continuous session risk score R_t at time t by aggregating normalized anomaly event weights:

$$R_t = \sum (w_i \cdot e_i(t)) + \lambda \cdot R_{t-1}$$

where $e_i(t) \in [0, 1]$ represents a specific detected event indicator, w_i denotes the empirical security weight assigned to that event category, and $\lambda \in [0, 1]$ represents an exponential memory decay factor that gradually diminishes past transient anomalies during prolonged periods of normal user activity.

The event indicators encompass: (1) e_{fail} : local proof witness calculation failure ($w_{fail} = 0.85$); (2) e_{geo} : sudden IP subnet or network origin shift ($w_{geo} = 0.50$); (3) e_{rate} : abnormal client request burst exceeding baseline thresholds ($w_{rate} = 0.40$); and (4) e_{wake} : browser tab resume following system sleep or background throttling ($w_{wake} = 0.35$).

TABLE I. ADAPTIVE RISK-TO-INTERVAL POLICY SPECIFICATION

State	Score Range & Indicators	Interval	Defensive Objective
Low	$R_t < 0.25$; steady interaction, no network shifts.	60 s	Minimize client CPU load and RPC bandwidth.
Medium	$0.25 \leq R_t < 0.60$; single anomaly (tab wake / rate jump).	30 s	Increase audit frequency during telemetry ambiguity.
High	$0.60 \leq R_t < 0.85$; multiple anomalies, repeated rate errors.	10 s	Compress vulnerability window to tight bounds.
Critical	$R_t \geq 0.85$; invalid proof, key mismatch, wallet disconnect.	Immediate	Instantaneous verification or local revocation (<100 ms).

C. Dynamic State Transition Mechanics

The scheduler operates as a finite state automaton governing transition between security levels: (1) Escalation Logic: A transition from LOW to MEDIUM occurs immediately when an individual anomalous event pushes R_t across the 0.25 threshold. If subsequent anomalies occur within the active window, the state escalates to HIGH ($R_t \geq 0.60$). A cryptographic pairing rejection or wallet provider severance immediately forces a transition to CRITICAL; (2) De-escalation Mechanics:

To prevent permanent lock-in from benign false positives, the system permits gradual de-escalation: if an elevated session completes a configurable consecutive run of successful verifications (e.g., three consecutive cycles) with no further anomalous events, the decay factor λ naturally attenuates R_t back to

LOW; (3) Irreversible Critical Rule: A transition into CRITICAL cannot be downgraded automatically; session tokens are purged immediately, requiring manual user re-authorization.

V. ADAPTIVE AUTHENTICATION ALGORITHM AND WORKFLOW

The operational execution flow proceeds according to the following formal algorithmic lifecycle:

1. Session Initialization:

User connects Web3 wallet provider (MetaMask); client retrieves public commitment H_{pub} and stores private secret s strictly in volatile WebAssembly heap memory closures. Risk state initializes to LOW ($I_0 = 60s$).

2. Continuous Telemetry Monitoring:

Client event listeners record user interaction velocity, tab visibility state changes (via HTML5 Page Visibility API), Web3 network events, and RPC response latencies.

3. Dynamic Score Evaluation:

At each evaluation interval or upon asynchronous interrupt, the risk engine updates $R_t = \Sigma (w_i \cdot e_i) + \lambda \cdot R_{t-1}$ and maps R_t to the corresponding interval $I_{next} \in \{60s, 30s, 10s, Immediate\}$.

4. Fresh Epoch Nonce Generation:

Prover queries local high-resolution clock and on-chain block timestamp to synthesize dynamic epoch freshness nonce n_t .

5. Client-Side Witness Proving:

snarkjs evaluates arithmetic Poseidon constraint Poseidon(s, n_t) = H_{pub} in WebAssembly, constructing Groth16 proof vector $argv = [\pi_A, \pi_B, \pi_C, \gamma_{pub}]$.

6. Asynchronous On-Chain Static Query:

ethers.js dispatches argv to Verifier.sol at 0x548d...5E6 on Polygon Amoy via static view call (eth_call), incurring exactly 0 wei in transactional gas fees.

7. Bilinear Pairing Verification:

Polygon Amoy RPC node evaluates pairing equations $e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(\gamma_{pub} \cdot K, \gamma) \cdot e(\pi_C, \delta)$ off-chain locally and returns boolean result.

8. Access Maintenance / Revocation:

If result = true, client session remains 'Verified: Access Granted', resetting fault counters. If result = false or timeout occurs, session tokens are expunged, and UI access drops in < 100 ms.

9. Interrupt-Driven Critical Handling:

If a wallet disconnects, account shift, or invalid private witness is detected out-of-band, the scheduler bypasses pending timers, immediately terminating session privileges.

10. Lifecycle Termination:

Execution loops continuously through steps 2–9 until explicit user logout, window closure, or security revocation.

VI. MATHEMATICAL FORMULATION OF CRYPTOGRAPHIC CIRCUITS

A. Poseidon Sponge Construction over BN254

To achieve sub-second client-side proving, our arithmetic circuit implements the Poseidon algebraic hash function over the prime scalar field F_p defined by the BN254 elliptic curve, where $p = 21888242871839275222246405745257275088548364400416034343698204186575808495617$. The circuit parameters are configured as $t = 2$ (state width of 2 elements, comprising secret witness s and dynamic nonce n_t), $R_F = 8$ full S-box rounds, and $R_P = 57$ partial rounds [1]. The round permutation applies non-linear power mappings $S(x) = x^5$ in F_p , which minimizes arithmetic gate complexity while providing robust resistance against differential, linear, and algebraic cryptanalysis.

The core circuit constraint verifies the pre-image relationship without revealing s :

$$\text{Poseidon}(s, n_t) = H_{\text{pub}}$$

where s is the private secret witness, n_t is the public epoch freshness nonce, and H_{pub} is the anchored public commitment.

B. Rank-1 Constraint System (R1CS) Representation

The Circom compiler formalizes the Poseidon permutation into a Rank-1 Constraint System (R1CS) satisfying:

$$(A \cdot z) \circ (B \cdot z) = (C \cdot z)$$

where $z = (1, \gamma_{\text{pub}}, w)$ is the complete witness vector containing public inputs $\gamma_{\text{pub}} = [H_{\text{pub}}, n_t]$ and

private intermediate wire assignments w , while A , B , and C are sparse coefficient matrices over F_p . By avoiding non-algebraic bitwise rotations, our Poseidon circuit requires exactly 240 R1CS constraints, compared to 25,120 constraints for SHA-256 and 148,300 constraints for Keccak-256 [1].

C. Groth16 zk-SNARK Protocol and Bilinear Pairing Checks

The Groth16 proving system transforms R1CS polynomial constraints into a succinct argument of knowledge comprising three group elements: $\pi_A \in G_1$, $\pi_B \in G_2$, and $\pi_C \in G_1$ over the pairing-friendly BN254 curve groups [1]. The on-chain Verifier contract verifies proof validity using the bilinear pairing mapping $e: G_1 \times G_2 \rightarrow G_T$:

$$e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(\sum \gamma_{\text{pub}} \{i\} \cdot K_i, \gamma) \cdot e(\pi_C, \delta)$$

where $\alpha \in G_1$, $\beta \in G_2$, $\gamma \in G_2$, $\delta \in G_2$, and $K_i \in G_1$ are immutable verification parameters derived during the trusted setup ceremony. If the prover constructed π using the authentic secret witness s , the pairing equality holds identically.

VII. COMPREHENSIVE THREAT MODELING AND SECURITY ANALYSIS

A. Post-Authentication Session Hijacking and Token Exfiltration

Under legacy Web2 and Web3 paradigms, exfiltrating an active session authorization token (e.g., via browser infostealer malware or XSS injection) grants the adversary indefinite access until token expiry. In our adaptive framework, an exfiltrated session token is mathematically decoupled from the private cryptographic witness s residing strictly in volatile WebAssembly memory. Because the Layer-2 Verifier contract demands continuous proof submissions bound to dynamic epoch freshness nonces n_t , an adversary possessing solely the bearer token cannot maintain access beyond the current verification interval (at most 10–60 seconds) without direct access to the client machine's protected memory sandbox [1], [4].

B. Proof Interception and Replay Attacks

If an adversary deploys a network packet sniffer between the client and the Polygon Amoy RPC endpoint, they may intercept the proof payload $\text{argv} = [\pi_A, \pi_B, \pi_C, \gamma_{\text{pub}}]$. However, because the public

signals γ_{pub} include the dynamic epoch freshness nonce $n_t = \text{timestamp} / \text{interval}$, the intercepted proof is valid exclusively for that specific temporal interval. When the attacker attempts to replay the proof in subsequent intervals, the Verifier contract calculates pairing checks against the current epoch nonce n_{t+1} , causing the pairing equation to fail and immediately invalidating the replayed proof [1], [15].

C. Man-in-the-Middle (MitM) and Calldata Tampering

If a hostile proxy alters any element of the proof vector argv in transit (e.g., attempting to substitute an attacker's public identity commitment H_{pub} for the legitimate user's commitment), the bilinear pairing equation $e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(\gamma_{\text{pub}} \cdot K, \gamma) \cdot e(\pi_C, \delta)$ fails mathematically. Furthermore, all RPC communication is encrypted using TLS 1.3, preventing payload inspection or tampering along the communication route [1], [5].

D. Denial of Service (DoS) and Adversarial Proving Floods

Because the adaptive scheduler accelerates verification upon detecting anomalous events, an attacker might deliberately inject forged network anomalies to exhaust the client's CPU or battery through repeated proof generation. Our framework counters this attack vector by enforcing an absolute lower bound on verification intervals ($I_{\text{min}} = 10\text{s}$), implementing strict sliding-window event rate limiters, and applying score decay for non-critical signals.

If anomaly triggers exceed configured rate thresholds, the system flags the anomalous traffic and drops the session rather than executing infinite proving loops [1], [14].

E. Host Environment and Operating System Compromise

If an adversary achieves root kernel-level access or deploys hardware keyloggers on the physical machine, any software-based authentication runtime can theoretically be compromised. To establish defense-in-depth against host-level compromises, the client-side proving architecture is engineered to integrate seamlessly with hardware-isolated enclaves, such as WebAuthn and FIDO2 cryptographic tokens. By delegating witness generation to a dedicated hardware secure enclave, private keys remain invulnerable even against compromised operating system kernels [1], [8].

VIII. IMPLEMENTATION DESIGN AND RUNTIME INTEGRATION

A. Client-Side WebAssembly Pipeline

The client runtime is engineered within a React single-page application integrating snarkjs and ethers.js [1]. The Poseidon cryptographic circuit is written in Circom and compiled to identity.wasm (1.2 MB) and identity_final.zkey (1.8 MB). Proving execution runs inside isolated Web Workers to prevent computational witness generation from blocking the browser's main UI thread. Peak memory consumption during Groth16 witness synthesis remains below 48 MB, ensuring compatibility with memory-constrained mobile browsers [1].

B. Layer-2 Smart Contract Verifier on Polygon Amoy

The Verifier.sol smart contract is deployed on the Polygon Amoy Layer-2 testnet at anchor address 0x548d484b506ce5Ef23CC8F4b34f20747A58ef5E6 [1]. The contract exports the public verifyProof(uint [2] _pA, uint[2][2] _pB, uint[2] _pC, uint[2] _pubSignals) view method. By executing verification calls as static eth_call queries, the EVM pairing precompile (0x08) evaluates the underlying elliptic curve operations off-chain locally within the RPC node, consuming zero on-chain gas fees (0 wei) while delivering verifiable cryptographic finality in sub-100 ms [1].

C. Tab Lifecycle and Suspension Management

Modern mobile and desktop web browsers enforce aggressive timer throttling on inactive background tabs, occasionally pausing JavaScript setInterval execution entirely. To ensure uninterrupted security continuity, our client integrates the HTML5 Page Visibility API: whenever a user switches away and subsequently returns to the application tab, the visibilitychange event immediately triggers an out-of-band re-verification cycle. This guarantees that an unattended session resuming after hours of sleep undergoes immediate cryptographic validation before unlocking sensitive application data [1].

IX. ANALYTICAL RESULTS AND SCHEDULE COMPARISON

Important scope note: The adaptive scheduling policy is evaluated here via analytical mathematical modeling across illustrative 60-minute session scenarios. These

calculations evaluate comparative scheduled proof counts and operational overhead; they do not represent measured benchmarks of a live production adaptive deployment.

For an observation window T , the scheduled proof count under a uniform interval I is $N \approx T / I$. For an adaptive schedule partitioned into k risk phases of duration T_i with corresponding intervals I_i , the cumulative scheduled proof count is $N_{\text{adapt}} = \sum_{i=1}^k (T_i / I_i)$.




Fig. 2. Analytical comparison of scheduled proof checks across 60-minute sessions under fixed and adaptive policies.

As illustrated in Fig. 2, an inflexible fixed 60-second schedule executes exactly 60 scheduled proof checks per hour. Under an illustrative Scenario 1 (50 minutes at Low risk and 10 minutes at medium risk), the adaptive schedule generates $50/1 + 10/0.5 = 70$ checks per hour. Under Scenario 2 (40 minutes at Low, 15 minutes at Medium, and 5 minutes at High risk), the schedule yields $40/1 + 15/0.5 + 5/0.166 = 100$ checks per hour. By comparison, aggressive fixed policies operating at 30-second and 10-second intervals require 120 and 360 checks per hour, respectively.

These findings illustrate that adaptive scheduling does not universally minimize proof count against every baseline; rather, it strategically allocates computational effort, enforcing rigorous 10-second audit density during suspect intervals while gracefully relaxing overhead to 60-second intervals during verified stable operation.

TABLE II. SCHEDULED CHECK COUNTS ACROSS 60-MINUTE SESSION SCENARIOS

Schedule Model	Operational Scenario Composition	Checks	Data (KB)
Fixed 60 s	60 min @ 60 s interval (Low risk baseline)	60	~72 KB

Adaptive (Scen. 1)	50 min @ 60s (Low) + 10 min @ 30s (Medium)	70	~84 KB
Adaptive (Scen. 2)	40 min @ 60s + 15 min @ 30s + 5 min @ 10s	100	~120 KB
Fixed 30 s	60 min @ 30 s interval (Aggressive baseline)	120	~144 KB
Fixed 10 s	60 min @ 10 s interval (High-security baseline)	360	~432 KB

X. EXPERIMENTAL METHODOLOGY FOR EMPIRICAL VALIDATION

To substantiate the theoretical models presented in Section IX, this section establishes a standardized, reproducible experimental methodology for empirical benchmarking across heterogeneous client and network environments. All evaluation runs must maintain identical hardware, browser engine, cryptographic circuit parameters, and Layer-2 contract addresses.

TABLE III. CONTROLLED EXPERIMENTAL EVALUATION PLAN

ID	Experiment Target	Primary Telemetry & Evaluation Metrics
E1	Fixed 60 s Baseline	Measure proof generation latency, CPU/RAM utilization, RPC load over 60 min.
E2	Adaptive Policy Loop	Record state transition accuracy, proof volume, and dynamic interval tracking.
E3	Benign Inactivity	Quantify false escalation rates and baseline power consumption over extended idle.
E4	Controlled Anomaly Injection	Inject synthetic network origin shift; measure time elapsed to interval tightening.
E5	Fault & Key Tampering	Supply incorrect secret witness; evaluate local circuit failure latency (<100 ms).

E6	RPC Disconnection & Failover	Sever primary RPC connection; evaluate fallback provider switch and graceful exit.
----	------------------------------	--

A. Multi-Platform Hardware Setup

Empirical testing will be executed across three physical device architectures: (1) High-Performance Workstation: Intel Core i7-13700H (14 cores, 20 threads), 16 GB DDR5 RAM, executing Chrome v124 (V8 WebAssembly engine); (2) ARM-Based Laptop: Apple M2 (8 cores), 16 GB unified memory, executing Safari v17.4 (Wasm JIT runtime); and (3) Mobile Device:

Qualcomm Snapdragon 8 Gen 2, 8 GB LPDDR5X RAM, executing Chrome Mobile v124 under Android 14. All devices connect via dedicated high-speed fiber Wi-Fi (802.11ax) to ensure uniform RPC transit times [1].

XI. EMPIRICAL BENCHMARKS AND COMPUTATIONAL PERFORMANCE

To establish baseline proving constraints, extensive experimental benchmarking was performed over $N = 100$ consecutive proof generation trials across heterogeneous hardware platforms running the optimized 240-constraint Poseidon circuit [1].

TABLE IV. MULTI-PLATFORM PROOF GENERATION BENCHMARKS ($N = 100$ TRIALS)

Hardware Architecture	Processor & Memory	Latency	Std. Dev.
Desktop Workstation	Intel Core i7-13700H, 16 GB	848.4 ms	± 34.2 ms
ARM-Based Laptop	Apple M2 (8-Core), 16 GB	712.1 ms	± 21.8 ms
Mobile Smartphone	Snapdragon 8 Gen 2, 8 GB	1,420.6 ms	± 88.5 ms

As shown in Table IV, desktop and laptop environments consistently complete witness generation and Groth16 proving in sub-second times (~ 712 ms to ~ 848 ms).

Even computationally restricted mobile smartphones complete proving within ~ 1.42 seconds, which fits easily inside the tightest 10-second high-risk heartbeat interval without stalling the user interface thread.

XII. COMPARATIVE DISCUSSION ACROSS DIGITAL IAM PARADIGMS

The empirical and analytical findings establish that synthesizing client-side zk-SNARK proving with adaptive Layer-2 verification addresses the fundamental trilemma of modern digital identity: security, user experience, and cost.

TABLE V. COMPARATIVE ANALYSIS ACROSS DIGITAL IAM PARADIGMS

Feature	Legacy Web2 (JWT)	Standard Web3 (SIWE)	Proposed Adaptive ZTA
Trust Model	Authenticate once, trust indefinitely.	Authenticate once, trust indefinitely.	Continuous Zero-Trust evaluation.
Witness Privacy	Credentials sent to central server.	Signed locally, visible on chain.	Remains strictly in client Wasm.
Risk Window	Hours to days until token expiry.	Indefinite until wallet disconnect.	Bounded to 10–60s epoch cutoff.
User Friction	Manual login / SMS MFA prompts.	Repetitive wallet signing pop-ups.	100% automated background proving.
Network Cost	Centralized database hosting bills.	High on-chain transaction gas fees.	0 wei (\$0.00 USD) via static eth_call.

As detailed in Table V, legacy OAuth/JWT architectures rely on perimeter defense, exposing users to hijacking when bearer tokens are harvested. Standard Web3 implementations (such as EIP-4361 Sign-In with Ethereum) authenticate wallet ownership initially, but revert to static session cookies for sustained operation to avoid transaction gas costs. Our framework is the first to achieve continuous, non-intrusive Zero-Trust authentication at exactly 0 wei in transactional gas fees, ensuring seamless user experience while eliminating the post-login vulnerability window [1], [6].

XIII. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

While the adaptive continuous authentication framework significantly advances Web3 session security, several operational considerations remain: (1) Risk Metric Calibration: Empirical weight assignments

(w_i) and score thresholds require machine learning calibration tailored to specific enterprise and dApp threat profiles; (2) Mobile Battery Considerations: While sub-second proving is feasible on flagship mobile devices, sustained 10-second proof cycles on low-end hardware can accelerate battery depletion; (3) Post-Quantum Cryptography (PQC): The Groth16 protocol relies on elliptic curve pairings over the BN254 curve, which remain vulnerable to future Shor's algorithm quantum cryptanalysis. Future research will explore lattice-based zero-knowledge proving primitives and recursive Incrementally Verifiable Computation (IVC) schemes (e.g., Nova and Halo2) to achieve quantum-resilient, cross-chain state folding [1], [14].

XIV. CONCLUSION

Static 'authenticate once, trust forever' architectures represent a fundamental flaw in modern digital identity systems, leaving decentralized applications vulnerable to post-login session hijacking and device takeover. This paper introduced an adaptive continuous authentication framework that combines browser-native recursive Zero-Knowledge Proofs with on-chain Layer-2 verification within a Zero-Trust architecture. By compiling an optimized 240-constraint Poseidon hashing circuit into WebAssembly, our system synthesizes cryptographic proofs in sub-second intervals (~712 ms to ~848 ms) while keeping private credentials isolated within client volatile memory closures.

Routing validation through read-only static view calls on Polygon Amoy completely eliminates on-chain gas fees (0 wei) while delivering sub-100 ms fault revocation upon credential failure or wallet disconnection. The adaptive scheduling engine strategically allocates verification frequency to session risk states, establishing a robust, privacy-preserving, and friction-free foundation for continuous identity management across modern decentralized ecosystems.

DECLARATIONS

Funding: No specific external funding was received for this work.

Conflict of Interest: The authors declare no competing financial or non-financial interests.

Data Availability: The schedule counts in Fig. 2 and Table II are derived analytical examples. Implementation source files are available upon reasonable request.

Code Availability: Circom arithmetic circuits, Solidity smart contracts, and React frontend integration scripts are available from the corresponding author upon reasonable request.

Author Contribution: Jayabandi Jersan developed the adaptive framework concept, cryptographic architecture, circuit designs, and authored the manuscript. Prof. Dhavala Lalitha Bhaskari supervised the research, validated methodology rigor, and reviewed the final manuscript.

REFERENCES

- [1] J. Jersan and D. L. Bhaskari, "A scalable continuous authentication system via recursive zero-knowledge proofs in a cross-chain zero-trust architecture," presented at QUEST-AI 2026, Andhra University, Visakhapatnam, India, 2026.
- [2] A. R. Raipurkar, S. Bobde, and A. Tripahi, "Digital identity system using blockchain-based self-sovereign identity & zero knowledge proof," in *Proc. 2023 OITS Int. Conf. Inf. Technol. (OCIT)*, Raipur, India, 2023, pp. 611–616, doi: 10.1109/OCIT59427.2023.10430981.
- [3] P. Zimmermann, "Why I wrote PGP," *PGP User's Guide*, 1991. [Online]. Available: Accessed: Aug. 14, 2023.
- [4] A. Whitten and J. D. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0," in *Proc. USENIX Security Symp.*, vol. 348, 1999, pp. 169–184.
- [5] N. Atzei, M. Bartoletti, and T. Cimoli, "A survey of attacks on Ethereum smart contracts (SoK)," in *Proc. 6th Int. Conf. Principles Security Trust (POST)*, vol. 10204, 2017, pp. 164–186.
- [6] Y. Yu, M. H. Au, G. Ateniese, X. Huang, W. Susilo, Y. Dai, and G. Min, "Identity-based remote data integrity checking with perfect data privacy preserving for cloud storage," *IEEE Trans. Inf. Forensics Security*, vol. 12, no. 4, pp. 767–778, 2016.

- [7] J. Shah, "What is blockchain verification & validation," 1Kosmos. [Online]. Available: Accessed: Aug. 17, 2023.
- [8] S. Jha, "The complete guide for types of blockchain!" Simplilearn. [Online]. Available: Accessed: May 31, 2023.
- [9] Bosch, "Digital identity – enabling secure collaboration with blockchain technology." [Online]. Available: Accessed: Jun. 30, 2023.
- [10] LeewayHertz, "What is zero knowledge proof and its role in blockchain." [Online]. Available: Accessed: Jul. 31, 2023.
- [11] Polygon Labs, "Polygon's zero knowledge strategy explained," Polygon. [Online]. Available: Accessed: Jan. 26, 2023.
- [12] M. de Vasconcelos Barros, F. Schardong, and R. Felipe Custódio, "Leveraging self-sovereign identity, blockchain, and zero-knowledge proof to build a privacy-preserving vaccination pass," *Social Science Research Network*, Feb. 2022, doi: 10.2139/ssrn.4036226.
- [13] E. Bandara, X. Liang, P. Foytik, S. Shetty, and K. De Zoysa, "A blockchain and self-sovereign identity empowered digital identity platform," in *Proc. Int. Conf. Comput. Commun. Netw. (ICCCN)*, 2021, pp. 1–7.
- [14] A. Grech, I. Sood, and L. Arino, "Blockchain, self-sovereign identity and digital credentials: Promise versus praxis in education," *Frontiers in Blockchain*, vol. 4, Art. no. 616779, 2021.
- [15] J. Hasan, "Overview and applications of zero knowledge proof (ZKP)," *International Journal of Computer Science and Network*, vol. 8, no. 5, pp. 436–440, 2019.
- [16] Q. Stokkink and J. Pouwelse, "Deployment of a blockchain-based self-sovereign identity," in *Proc. 2018 IEEE Int. Conf. Internet Things (iThings) and IEEE GreenCom, CPSCom and SmartData*, 2019, pp. 1336–1342.
- [17] X. Liang, S. Shetty, J. Zhao, D. Bowden, D. Li, and J. Liu, "Towards decentralized accountability and self-sovereignty in healthcare systems," in *Proc. Int. Conf. Inf. Commun. Security (ICICS)*, vol. 10631, 2018, pp. 387–398.