MEDVIX PUBLICATIONS

Journal of Medicine Care and Health Review
Research Article | Open Access | ISSN 3065-1719 | Volume 03 | Issue 01

Blockchain-Anchored Adaptive Authentication for Real-Time Medical Data Streams in AI-Driven Smart Grid-IoMT Converged Networks

Sanaz Farhadi[1*], Uras Panahi[2]

[1]Shiraz University, School of Biomedical Engineering, Shiraz, Iran.

[2]Sakarya University, Department of Computer Engineering, Serdivan, Sakarya, Turkey.

Received date: 08 June 2026; **Accepted date:** 20 June 2026; **Published date:** 26 June 2026

Corresponding Author: Sanaz Farhadi, Shiraz University, School of Biomedical Engineering, Shiraz, Iran.

Citation: Sanaz Farhadi, Uras Panahi, Blockchain-Anchored Adaptive Authentication for Real-Time Medical Data Streams in AI-Driven Smart Grid-IoMT Converged Networks. Journal of Medicine Care and Health Review 3(1). https://doi.org/10.61615/JMCHR/2026/JUNE027140626

Abstract

The convergence of Smart Grids and the Internet of Medical Things (IoMT), termed Grid-IoMT, represents an emerging paradigm where healthcare facilities dynamically interact with energy grids to optimize both clinical operations and power consumption. Real-time medical data streams (e.g., continuous vital signs from wearable monitors, infusion pump logs, ventilatory parameters) traverse network infrastructure shared with grid telemetry, creating unprecedented attack surfaces where energy-demand manipulation can indirectly compromise patient safety, and conversely, medical data injection can destabilize grid frequency regulation. This paper presents BlockAuth-GridMed, a novel blockchain-anchored adaptive authentication framework specifically designed for real-time medical data streams in AI-driven Smart Grid-IoMT converged networks. The framework integrates three synergistic innovations:

(1) A hierarchical blockchain architecture (local permissioned chains for clinical domains interconnected via a main chain for cross-domain trust) that anchors authentication proofs without introducing latency prohibitive for real-time medical applications (median latency 187ms),

(2) An adaptive authentication engine powered by deep reinforcement learning (DRL) that dynamically adjusts authentication strength based on real-time risk assessment—escalating to multi-factor requirements during grid instability events or cyber-threat alerts while maintaining low-friction single-factor authentication during quiescent periods,

(3) A zero-knowledge proof (ZKP) layer enabling mutual authentication between medical devices and grid nodes without revealing sensitive patient identifiers or clinical data patterns to energy system operators.

We evaluate BlockAuth-GridMed using a realistic testbed emulating a 200-bed smart hospital integrated with an IEEE 13-bus distribution grid model, processing 15,000 real-time medical data streams per second across 6,500 IoMT devices and 12 grid sensors. The framework achieves 99.97% authentication success rate for latency-sensitive medical alerts (critical events requiring <100ms end-to-end latency) and maintains an average authentication overhead of 28ms, well within clinical requirements. Under adversarial conditions (simulated man-in-the-middle, replay, and grid-state injection attacks), BlockAuth-GridMed demonstrates 96.8% attack detection and 99.1% attack prevention rates, outperforming baseline certificate-based (83.4%/87.2%) and token-based (71.3%/74.6%) schemes. The DRL-driven adaptive authentication reduces unnecessary multi-factor challenges by 73% compared to static high-security policies, significantly improving clinical workflow efficiency. We also analyze blockchain gas costs (approx. \$0.012 per authentication), scalability under IoMT device churn (up to 15% daily device joins/leaves), and regulatory alignment with HIPAA, NERC CIP, and FDA pre-market guidance for medical device security. This work provides the first integrated authentication framework specifically tailored to the Grid-IoMT convergence, enabling secure, real-time, and adaptive protection for medical data streams in energy-aware healthcare infrastructures.

Keywords: Blockchain, adaptive authentication, Internet of Medical Things (IoMT), Smart Grid, real-time medical data streams, zero-knowledge proofs, deep reinforcement learning, cyber-physical security, converged networks.

Introduction

The Grid-IoMT Convergence

Healthcare facilities are among the most energy-intensive commercial buildings, with a typical 200-bed hospital consuming 2.5–3.5 MW of average power and experiencing demand peaks exceeding 5 MW. Simultaneously, the proliferation of IoMT devices from wearable cardiac monitors to smart infusion pumps has created data-intensive clinical environments where real-time physiological streams demand guaranteed low-latency network performance. Historically, hospital energy management and clinical IT networks operated as disjoint domains: facilities management controlled HVAC, lighting, and medical equipment power; clinical engineering managed IoMT devices and data networks. However, three converging trends are forcing integration:

Demand Response Participation: Hospitals increasingly participate in demand response (DR) programs where grid operators request load reduction

during peak conditions, receiving financial incentives. Real-time adjustments to non-critical medical equipment (e.g., HVAC in non-patient areas, battery charging schedules for mobile IoMT devices) require bidirectional communication between grid operators and hospital energy management systems networks that also carry time-sensitive medical data **[1]**.

AI-Driven Grid Optimization: Grid operators deploy AI models that forecast hospital energy demand using patterns derived from operational data. Without proper isolation, these models might inadvertently infer patient census, procedure schedules, or even individual patient acuity from aggregated power signatures, raising privacy concerns that demand cryptographic protection.

Edge Computing for Clinical AI: AI models processing real-time medical data (e.g., sepsis prediction from vital signs) are increasingly deployed at edge gateways within hospital networks. These same gateways often aggregate grid telemetry, creating convergence points where medical and energy data streams coexist.

Unique Security Challenges of Grid-IoMT Convergence

Converged Grid-IoMT networks introduce authentication requirements not addressed by existing solutions

Challenge 1: Real-Time Latency Constraints – Medical alerts (e.g., asystole detection from ECG monitors) require end-to-end latency <100ms from sensor to clinical decision support. Blockchain-based authentication systems, with typical latencies of seconds for consensus, are unusable for such streams without novel architectures **[2]**.

Challenge 2: Cross-Domain Trust Without Full Visibility – Grid operators must authenticate that a load reduction request originates from an authorized hospital device, but they must not learn which patient or clinical condition triggered the request (privacy). Conversely, clinical systems must authenticate grid state information (e.g., "imminent brownout") without granting grid operators access to patient data.

Challenge 3: Variable Risk Environment – Authentication requirements should be context-aware: during grid instability (e.g., frequency deviations suggesting cyberattack), authentication should escalate to multi-factor; during stable periods, low-friction authentication preserves clinical workflow efficiency.

Challenge 4: IoMT Device Heterogeneity and Churn – IoMT environments include devices with vastly different capabilities (from implantable sensors with μW power budgets to bedside monitors with ample compute). Authentication mechanisms must accommodate this spectrum.

Limitations of Existing Approaches

Approach Limitation for Grid-IoMT

PKI/certificates Centralized certificate authorities create single points of failure; revocation is slow.

OAuth2 / tokens assumes continuous connectivity to authorization servers; it fails during grid isolation events.

Biometric authentication Not feasible for machine-to-machine (M2M) IoMT-grid interactions.

Traditional blockchain (Ethereum) Consensus latency (seconds to minutes) violates real-time medical requirements.

Physical unclonable functions (PUFs) are device-specific but lack revocation mechanisms and cross-domain trust **[3]**.

Contributions

This paper presents BlockAuth-GridMed, the first blockchain-anchored adaptive authentication framework specifically designed for real-time medical data streams in Grid-IoMT converged networks. Contributions are:

1. Hierarchical blockchain architecture with local permissioned chains for clinical subdomains (latency <200ms) interconnected via a main chain for cross-domain notarization.
2. Deep reinforcement learning (DRL) adaptive authentication engine that dynamically selects authentication factors (single-factor, two-factor, biometric, or hardware token) based on real-time risk scores from grid state, network threat intelligence, and clinical context.
3. Zero-knowledge proof (ZKP) layer enabling mutual authentication across clinical and grid domains without revealing sensitive identifiers.
4. Comprehensive evaluation on a realistic testbed (200-bed hospital + IEEE 13-bus grid) under normal and adversarial conditions.
5. Regulatory and economic analysis, including HIPAA, NERC CIP alignment, and gas cost modelling **[4]**.

Background and Related Work

Smart Grid-IoMT Convergence Architectures

Recent work has explored technical architectures for healthcare-energy integration. The ENERGY-MED project (2024-2026) demonstrated that hospitals can reduce peak demand by 18-24% through real-time scheduling of non-critical medical device charging without impacting patient care. However, their security analysis assumed isolated network domains and did not address cross-domain authentication. Similarly, Alshehri and colleagues proposed a middleware layer for hospital-grid integration but relied on pre-shared keys unsustainable at the IoMT scale (thousands of devices).

Authentication for Medical Data Streams

Existing IoMT authentication schemes fall into three categories:

(1) Lightweight cryptographic protocols (e.g., ECC-based handshakes) that assume relatively stable network environments;

(2) Physically unclonable function (PUF)-based authentication suitable for resource-constrained devices but lacking revocation;

(3) Blockchain-based systems that anchor device identities but generally assume latency-tolerant applications. A 2025 systematic review by Mahmood et al. identified that no existing solution addresses the Grid-IoMT convergence specifically[5].

Adaptive Authentication

Adaptive (risk-based) authentication adjusts requirements based on contextual risk. Widely deployed in consumer finance (e.g., step-up authentication for high-value transactions), adaptive methods have been proposed for healthcare. However, existing healthcare adaptive authentication considers only clinical context (patient acuity, data sensitivity), not external grid state or cyber-threat intelligence, both critical in Grid-IoMT.

Blockchain for Cyber-Physical Systems

Permissioned blockchains (Hyperledger Fabric, Corda) have been deployed for industrial IoT authentication, achieving latencies of 100-500ms approaching real-time feasibility. However, existing designs assume homogeneous device capabilities and do not accommodate the extreme heterogeneity of IoMT devices, from implantables to imaging systems.

Research Gap Identified

No prior work provides: (1) integrated authentication spanning clinical and grid domains with formal privacy guarantees; (2) adaptive mechanisms informed by both clinical risk and grid state; (3) latency performance meeting real-time medical alert requirements (<100ms for critical alerts); (4) support for the full spectrum of IoMT device capabilities **[6]**.

BlockAuth-GridMed Framework

System Architecture Overview

BlockAuth-GridMed adopts a four-layer architecture:

Layer 1: IoMT Device Layer – Medical devices (wearable monitors, infusion pumps, ventilators, implantable sensors) and grid sensors (smart meters, PMUs). Devices range from Class 0 (ultra-constrained: implantables, 8KB RAM) to Class 2 (bedside monitors, 512MB+ RAM).

Layer 2: Edge Authentication Layer – Hospital edge gateways and grid edge nodes that perform local authentication caching, risk scoring, and act as blockchain clients.

Layer 3: Hierarchical Blockchain Layer – Local permissioned chains (Hyperledger Fabric) deployed per clinical unit (ICU, ED, OR) and grid subnet, plus a main chain anchoring cross-domain proofs.

Layer 4: AI Risk Assessment Layer – DRL agents that compute real-time authentication risk scores using inputs from grid state monitors, threat intelligence feeds, and clinical context.

Hierarchical Blockchain for Low-Latency Authentication

Local Chains: Each clinical unit operates a permissioned Fabric chain with 3-5 ordering nodes. Consensus uses Raft (crash fault-tolerant, not Byzantine), sufficient for a hospital's internal trust model where all nodes are administered by a single organization. Local chain latency: median 87ms, 95th percentile 142ms **[7]**.

Main Chain: The inter-domain chain uses a more robust consensus (PBFT, 7 nodes spanning hospital and grid operator organizations). Main chain latency: median 412ms, but it is only invoked for cross-domain authentication (e.g., a grid node authenticating a medical device's load reduction request). For intra-hospital authentication, local chains suffice.

Authentication Transaction Structure: Each transaction contains: device ID (hashed), timestamp, authentication proof (digital signature or ZKP), validity window, and action type (data publish, control command, grid request).

Caching and Pre-validation: Edge gateways cache recent authentication outcomes. For recurring device-grid interactions (e.g., same device sending periodic vital sign streams), the gateway validates against cached state and only anchors to blockchain every Nth transaction (N=50 for high-frequency devices), reducing blockchain load by 98%.

Deep Reinforcement Learning for Adaptive Authentication

State Space: The DRL agent observes:

- Grid risk score (0-1): based on frequency deviation, voltage instability, and active cyber alerts from grid IDS
- Network threat score (0-1): based on recent anomalous traffic patterns in the hospital network
- Clinical context: patient acuity (ICU vs. general ward), data sensitivity (PHI vs. anonymous statistics)
- Device attributes: class, authentication history (failed attempts), battery level
- Time context: time of day, day of week (attack patterns vary)

Action Space: Authentication factor requirements:

- Action A0: Single-factor (device ID + timestamp)
- Action A1: Two-factor (ID + time-based one-time password)
- Action A2: Hardware token challenge (for high-risk)
- Action A3: Biometric (user-facing devices only)
- Action A4: Full re-enrollment (blockchain re-anchoring)

Reward Function: $R = \alpha \times Security\_Score - \beta \times User\_Friction\_Cost - \gamma \times Latency\_Penalty$

Security Score: 1.0 if authentication resists simulated attack, decreasing with vulnerabilities.

User Friction Cost: 0 for single-factor, 0.3 for 2FA, 0.7 for hardware token, 0.9 for biometric.

Latency Penalty: penalty proportional to authentication latency exceeding device-specific threshold **[8]**.

Algorithm: Proximal Policy Optimization (PPO) with 3-layer MLP policy network (256-128-64). Training on 6 months of simulated Grid-IoMT activity (2.1 million authentication events). Converged after 45,000 episodes **[9]**.

Zero-Knowledge Proof Layer for Cross-Domain Privacy

Clinical devices must prove authorization to grid operators (e.g., "this load reduction request is legitimate") without revealing which patient or clinical condition triggered it. BlockAuth-GridMed implements ZK-SNARKs (Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge) with the following statement:

"I (device with identity commitment C) am authorized to send message M of type T (where $T \in$ {load_reduction, grid_state_query, clinical_data}) under policy P, and I possess a valid authentication token anchored at blockchain block B."

Proof Generation: Performed by edge gateways (Class 2 devices lack compute for ZK proofs). Proof size: 288 bytes; verification time: 12ms; proof generation: 340ms (acceptable for cross-domain interactions that are less frequent than intra-domain).

Privacy Guarantee: Grid operators verify the proof without learning device identity, patient context, or clinical data content.

Registration and Revocation Protocol

Device Registration: Each IoMT device generates a key pair during commissioning. The public key and device class are recorded on the local blockchain with a validity period. A zero-knowledge commitment to device identity is generated for cross-domain use.

Revocation: If a device is compromised, the hospital administrator issues a revocation transaction on the local chain. Edge gateways receive this within

the local chain's block interval (2 seconds). For cross-domain, the main chain records a revocation proof grid node; check this before accepting cross-domain authentications.

Re-authentication Period: Devices must re-anchor every 24 hours (for continuous streaming) or per session for intermittent interactions.

Resource-Optimized Variants

Class 0 (Implantable sensors, <32KB RAM): Cannot run blockchain client or perform asymmetric crypto per transaction. Instead: (1) device maintains a pre-shared symmetric key established during registration; (2) edge gateway performs authentication on the device's behalf using cached state; (3) device only transmits a lightweight MAC (message authentication code) per packet. Security: weaker but acceptable for low-risk data types (e.g., temperature monitoring, not infusion control) **[10]**.

Class 1 (Wearables, 32-256KB RAM): Support ECC-256 signatures but not a full blockchain client. Offload blockchain interactions to edge gateway; device signs each transmission; gateway verifies and anchors batch signatures.

Class 2 (Bedside monitors, >256KB RAM): Run a lightweight blockchain client (simplified payment verification mode) and participate in the local chain as a light node **[3]**.

Experimental Methodology

Testbed Description

We constructed a Grid-IoMT testbed emulating a 200-bed smart hospital (modeled on University of Pittsburgh Medical Center data) integrated with an IEEE 13-bus distribution grid model.

IoMT Devices

- 4,500 wearable patient monitors (Class 1, 2/minute vital signs)
- 1,200 smart infusion pumps (Class 1/2, event-driven)
- 600 bedside monitors (Class 2, continuous)
- 200 implantable cardiac devices (Class 0, periodic)
- Total: 6,500 devices generating 15,000 data streams/second (simulated)

Grid Components

- 12 smart meters at hospital load centers
- 2 PMUs (phasor measurement units)
- Grid edge node running DRL risk assessment

Network: Isolated 10 GbE backbone with simulated WAN latency (5-25ms) between hospital and grid operator.

Blockchain

- 15 local chains (one per clinical unit)
- 1 main chain (7 nodes: 5 hospitals, 2 grid operators)

Baseline Systems Compared

Baseline Description

- Baseline 1: Certificate-based (X.509) Traditional PKI with short-lived certificates (1 hour validity)
- Baseline 2: Token-based (OAuth2) Centralized authorization server issuing JWTs
- Baseline 3: Static blockchain (no adaptation) Same blockchain architecture but fixed authentication policy (always 2FA)
- Baseline 4: PUF-based Physical unclonable function per device, no revocation

Adversarial Scenarios

Scenario A (MITM): Adversary intercepts the authentication handshake and attempts to replay or modify.

Scenario B (Grid state injection): Adversary injects false frequency readings to trigger inappropriate authentication escalation.

Scenario C (Device cloning): Adversary captures device credentials and attempts to impersonate.

Scenario D (Replay attack): Adversary replays captured authentication messages.

Scenario E (Cross-domain poisoning): Adversary compromises the grid node and attempts to impersonate a clinical device.

Evaluation Metrics

- Authentication success rate (target >99.9%)
- Authentication latency (median, 95th, 99th percentile)
- Attack detection rate (true positive)
- Attack prevention rate (attacks blocked)
- False positive rate (legitimate authentication incorrectly challenged)
- Reduction in unnecessary multi-factor challenges
- Blockchain gas cost per authentication
- Device churn tolerance (new devices joining/leaving network)

Results

Authentication Performance Under Normal Conditions

BlockAuth-GridMed achieved 99.97% authentication success rate across all medical data streams. Critical alerts (<100ms end-to-end requirement): 98.2% satisfied within 80ms; 99.4% within 100ms.

Table 1 — Authentication Latency by Device Class and Authentication Type

Performance across 6,500 IoMT devices under normal operating conditions

Device Class	Authentication Type	Median (ms)	95th %ile (ms)	Success Rate
Class 0 (implantable)	Edge-cached MAC	14	32	99.4%
Class 1 (wearable)	Single-factor	23	47	99.9%
Class 1 (wearable)	2FA (DRL-triggered)	58	89	99.9%
Class 2 (bedside)	Single-factor	18	41	99.98%
Class 2 (bedside)	Hardware token	87	134	99.99%
Cross-domain	ZKP + main chain	412	689	99.92%

Citation: Sanaz Farhadi, Uras Panahi, Blockchain-Anchored Adaptive Authentication for Real-Time Medical Data Streams in AI-Driven Smart Grid-IoMT Converged Networks. Journal of Medicine Care and Health Review 3(1).
https://doi.org/10.61615/JMCHR/2026/JUNE027140626

Observation: Even worst-case cross-domain ZKP authentication (689ms 95th percentile) meets non-critical medical data requirements (e.g., hourly reporting) but exceeds critical alert latency bounds. Critical alerts use pre-established sessions with cached authentication.

Adaptive Authentication Effectiveness

The DRL agent reduced unnecessary multi-factor challenges by 73% compared to a static high-security policy (always 2FA), while maintaining equivalent security outcomes **[6]**.

Table 2 — Adaptive Authentication Impact on Clinical Workflow

DRL-driven adaptation vs. static security policies (daily averages)

Security Policy	Multi-Factor Challenges (daily)	Authentication Failures (patient-reported friction)	Security Score (0-1)
Static low (always single-factor)	0	2.1	0.71
Static medium (always 2FA)	12,400	287	0.94
Static high (hardware token always)	12,400	412	0.97
BlockAuth-GridMed (DRL-adaptive)	**3,348**	**79**	**0.93**

Key finding: 73% reduction in unnecessary 2FA challenges directly reduced clinical workflow disruption (287 patient-reported friction events to 79).

Risk Score Drivers: Feature importance analysis from DRL policy revealed grid state (34% weight) and network threat score (28%) as top drivers for authentication escalation, confirming that Grid-IoMT convergence fundamentally changes authentication requirements compared to isolated medical networks.

Security Against Adversarial Attacks

Table 3 — Attack Detection and Prevention Rates (5 Adversarial Scenarios)

Comparison with certificate-based, token-based, and static blockchain baselines

Attack Scenario	BlockAuth-GridMed Detection	BlockAuth-GridMed Prevention	Baseline 1 (Certificate)	Baseline 2 (Token)	Baseline 3 (Static BC)
MITM (A)	97.2%	99.4%	88.1% / 91.2%	74.3% / 78.1%	95.1% / 97.3%
Grid state injection (B)	94.1%	98.2%	76.2% / 79.4%	68.2% / 71.3%	86.4% / 90.2%
Device cloning (C)	98.4%	99.7%	85.4% / 89.1%	71.2% / 74.0%	96.3% / 98.1%
Replay (D)	96.8%	99.3%	82.3% / 86.7%	70.1% / 73.2%	94.2% / 96.8%
Cross-domain poisoning (E)	97.5%	99.1%	84.8% / 87.2%	72.9% / 75.8%	93.8% / 96.1%
Average	**96.8%**	**99.1%**	**83.4% / 87.2%**	**71.3% / 74.6%**	**93.2% / 95.7%**

Interpretation: BlockAuth-GridMed outperforms certificate and token baselines across all scenarios. The static blockchain baseline (fixed authentication policy, no adaptation) achieves strong prevention (95.7% average) but imposes high friction (always 2FA). BlockAuth-GridMed matches this security level while reducing friction by 73%.

False Positive Rate: Legitimate authentications incorrectly flagged as suspicious: 1.2% for BlockAuth-GridMed, 2.8% for static blockchain (higher because fixed policies cannot reduce suspicion when context indicates safety).

Blockchain Performance and Scalability

Table 4 — Blockchain Consensus Latency and Throughput

Local chains (Raft) vs. Main chain (PBFT) performance

Chain Type	Consensus	Nodes	Median Latency	Throughput (tx/sec)	Max Capacity (devices)
Local (clinical unit)	Raft	3-5	87ms	850	1,200
Local (grid subnet)	Raft	3	92ms	720	500
Main chain	PBFT	7	412ms	280	N/A (cross-domain only)

Citation: Sanaz Farhadi, Uras Panahi, Blockchain-Anchored Adaptive Authentication for Real-Time Medical Data Streams in AI-Driven Smart Grid-IoMT Converged Networks. Journal of Medicine Care and Health Review 3(1).
https://doi.org/10.61615/JMCHR/2026/JUNE027140626

Gas cost: Per authentication transaction on local chain: $0.008–0.012 (assuming Hyperledger Fabric on private cloud; public chain equivalent would be ~$0.12, prohibitive). Annual cost for 6,500 devices authenticating at 1/hour average: ~$1,800, acceptable for hospital budgets.

Device churn tolerance: When 15% of devices join or leave daily (simulating patient turnover), local chain reconfiguration takes 18 seconds on average, sufficiently fast that authentication caches remain valid during transition **[2]**.

Resource Consumption on IoMT Devices

Table 5 — Device Resource Overhead per Authentication

Impact on RAM, energy, and battery life across device classes

Device Class	Operation	RAM (KB)	Energy (μJ)	Battery Impact (daily %)
Class 0 (implantable)	MAC generation	4	0.8	0.02%
Class 1 (wearable)	ECC-256 sign	28	142	0.4%
Class 1 (wearable)	2FA (TOTP)	18	34	0.1%
Class 2 (bedside)	Full blockchain light client	1,024	2,400	2.1%

Observation: Even Class 0 devices (implantables) can support lightweight MAC authentication with negligible battery impact (<0.1% daily). Full blockchain participation is restricted to Class 2 devices plugged into line power (bedside monitors).

Discussion

Key Findings

1. **Feasibility of blockchain for real-time medical authentication:** With hierarchical architecture (local + main chain) and aggressive caching, median latency of 87ms meets requirements for all but the most latency-critical alerts (which use pre-established sessions).
2. **DRL adaptation effectively balances security and friction:** 73% reduction in unnecessary 2FA challenges while maintaining 0.93 security score (vs. 0.94 for static 2FA).
3. **Zero-knowledge proofs enable cross-domain privacy:** Grid operators authenticate medical device requests without learning patient identifiers essential for regulatory compliance.
4. **Grid state is a first-class authentication factor:** Grid instability accounted for 34% of risk score variance; ignoring this factor would leave converged networks vulnerable **[5]**.

Limitations

Byzantine fault tolerance assumptions: Local chains use Raft (crash fault-tolerant), assuming hospital network nodes are trusted. A compromised hospital administrator node could corrupt the local chain. For higher assurance, migration to Byzantine-fault-tolerant (BFT) consensus (e.g., Tendermint) is possible but increases latency to ~400ms.

ZK proof generation overhead: 340ms generation time by edge gateways is acceptable for cross-domain interactions (which are less frequent), but cannot be performed per-packet. Mitigation: Establish long-lived (24-hour) ZK credentials renewed daily.

Implantable device security: Class 0 devices use symmetric MACs, which are vulnerable if the edge gateway is compromised. In critical applications (e.g., pacemakers requiring authentication for programming), a hardware trust anchor (secure element) should be added, though this increases cost and energy consumption.

Real-world deployment complexity: Integrating with legacy hospital systems (some predating modern security practices) requires middleware adapters. Our testbed assumed clean-slate integration; real-world deployment would face substantial engineering challenges.

Regulatory Alignment

HIPAA: BlockAuth-GridMed's zero-knowledge layer ensures that grid operators never receive PHI. Local blockchain transactions contain hashed device IDs (not direct identifiers) and encrypted clinical data pointers compliant with HIPAA de-identification standards **[9]**.

NERC CIP (grid security): The framework provides non-repudiation for demand response commands, a NERC CIP requirement for critical cyber assets. Grid operators can verify that load reduction requests originate from authorized hospital devices without learning clinical details.

FDA pre-market guidance: For medical devices implementing BlockAuth-GridMed, the authentication mechanism would likely be classified as a "non-critical software function" (assuming it does not directly control therapy delivery). However, devices that rely on authentication for safety-critical commands (e.g., infusion pump dose changes) would require FDA pre-market review.

Future Directions

Post-quantum cryptography: Current ECC-256 signatures are vulnerable to quantum attacks (expected 2030-2035). Transition to CRYSTALS-Dilithium or SPHINCS+ for blockchain transactions is planned.

Decentralized identity (DID) integration: Replace device IDs with W3C-compliant DIDs, enabling cross-organizational interoperability without a central registration authority.

Federated learning for risk assessment: Instead of centralized DRL training, distribute risk model training across hospitals using federated learning, improving generalization while preserving data privacy.

Hardware acceleration: For Class 0 devices, integrate lightweight hardware security modules (HSMs) supporting ECC-256 with ultra-low power (existing commercial modules: <5 μJ per signature).

Conclusion

Citation: Sanaz Farhadi, Uras Panahi, Blockchain-Anchored Adaptive Authentication for Real-Time Medical Data Streams in AI-Driven Smart Grid-IoMT Converged Networks. Journal of Medicine Care and Health Review 3(1). https://doi.org/10.61615/JMCHR/2026/JUNE027140626

The convergence of Smart Grids and IoMT networks creates a new security paradigm where medical data streams traverse infrastructure shared with energy telemetry. Existing authentication mechanisms designed for isolated clinical or grid domains fail to address the unique requirements of Grid-IoMT: real-time latency constraints, cross-domain privacy, variable risk environments, and extreme device heterogeneity.

BlockAuth-GridMed provides the first integrated solution tailored to this convergence. Through hierarchical blockchain architecture (achieving 87ms median authentication latency), DRL-driven adaptive authentication (73% reduction in unnecessary multi-factor challenges while maintaining 0.93 security score), and a zero-knowledge proof layer for cross-domain privacy, the framework enables secure, privacy-preserving authentication for real-time medical data streams in energy-aware healthcare infrastructures.

Evaluation on a realistic 200-bed hospital testbed with 6,500 IoMT devices and an integrated grid model demonstrated 99.97% authentication success rate, 96.8% attack detection, and 99.1% attack prevention, substantially outperforming certificate and token baselines. The framework aligns with HIPAA, NERC CIP, and FDA guidance, providing a regulatory-ready path for deployment.

As healthcare facilities increasingly participate in demand response and grid operators incorporate hospital load forecasting into AI-driven optimization, the Grid-IoMT convergence will accelerate. BlockAuth-GridMed offers a foundational authentication layer that enables this convergence without compromising patient safety, privacy, or clinical workflow efficiency. Future work will address post-quantum security, decentralized identity integration, and real-world clinical pilot deployments.

References

1. Umran, S. M., Lu, S., Abduljabbar, Z. A., & Tang, X. (2023). A blockchain-based architecture for securing industrial IoTs data in electric smart grid. Computers, Materials & Continua, 74(3), 5389–5416.
2. Zahoor, A., Mahmood, K., Shamshad, S., Saleem, M. A., Ayub, M. F., Conti, M., & Das, A. K. (2023). An access control scheme in IoT-enabled smart-grid systems using blockchain and PUF. Internet of Things, 22, 100708.
3. Saleem, M. A., Li, X., Mahmood, K., Shamshad, S., Ayub, M. F., Bashir, A. K., & Omar, M. (2024). Provably secure conditional-privacy access control protocol for intelligent customers-centric communication in VANET. IEEE Transactions on Consumer Electronics, 70(1), 1747–1756.
4. Hao, X., et al. (2023). A blockchain-based cross-domain and autonomous access control scheme for Internet of Things. IEEE Transactions on Services Computing, 16(2), 1–14.
5. Farooq, M. S., et al. (2026). Fuzzychain-edge: Adaptive blockchain access control using fuzzy logic and zero-knowledge proofs for IoT-edge healthcare. IEEE Access, 13, 18660–18676.
6. Li, D., Yang, Z., & Yu, S. (2024). A micro-segmentation method based on VLAN-VxLAN mapping technology. Future Internet, 16(9), 320.
7. Al Hwaitat, A. K., Almaiah, M. A., Ali, A. (2023). A new blockchain-based authentication framework for secure IoT networks. Electronics, 12(17), 3618.
8. Li, Z., Li, K., et al. (2026). A lightweight mutual authentication and key exchange protocol for resource-constrained IoMT devices. In International Conference on Network and System Security (NSS 2025), 133–145. Springer.
9. Alnahari, W., & Quasim, M. T. (2021). Privacy concerns, IoT devices, and attacks in smart cities. In 2021, the International Congress of Advanced Technology and Engineering (ICOTEN) 1–5. IEEE.
10. Mahmood, K., Obaidat, M. S., Shamshad, S., Alenazi, M. J. F., Kumar, G., Anisi, M. H., & Conti, M. (In press). Cost-effective authenticated solution (CAS) for 6G-enabled artificial intelligence of medical things (AIoMT). IEEE Internet of Things Journal.