Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 2, Issue 1, January – March 2025 || PP. 31-40 https://doi.org/10.63345/sjaibt.v2.i1.104

AI-Blockchain Integration for Remote Patient Monitoring and Alerts

Dr. Jürgen Schneider
Faculty of Quantum Studies
Zurich Academy of Science, Switzerland

Date of Submission: 27-12-2024 **Date of Acceptance: 31-12-2024** **Date of Publication: 01-01-2025**

ABSTRACT

Remote Patient Monitoring (RPM) has matured from episodic teleconsultations to continuous, sensor-driven care supported by edge analytics and cloud services. Yet three friction points persist: (i) privacy and trust in data handling, (ii) interoperability across fragmented health information systems, and (iii) timely, auditable alerting that can be verified across organizations. This manuscript proposes a reference architecture that fuses Artificial Intelligence (AI) for streaming physiological inference with permissioned blockchain for tamper-evident logging, consent management, and cross-institutional data exchange. The design: (1) captures multi-modal signals from Internet of Medical Things (IoMT) wearables; (2) performs on-device/edge AI for anomaly detection to minimize latency and exposure of raw data; (3) persists summaries and cryptographic hashes on a consortium blockchain while storing large payloads off-chain; (4) implements smart-contract-based consent and alert workflows integrated with HL7® FHIR® resources; and (5) supports privacy-preserving learning (federated learning) to continually improve models without centralizing protected health information (PHI). We synthesize the most recent guidance from WHO and FDA on telemedicine/RPM, FHIR-based interoperability, and global privacy frameworks (HIPAA, GDPR, India's DPDP Act), and we translate these into concrete design controls. A qualitative evaluation across security, latency, and governance dimensions suggests the approach can reduce alert time, strengthen provenance and auditability, and enable compliant data minimization. We conclude with an implementation roadmap, known risks (e.g., key management, model drift, edge heterogeneity), and future research directions such as zero-knowledge proofs for selective disclosure and formal verification of smart-contracted clinical alerts.

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 2, Issue 1, January – March 2025 || PP. 31-40 https://doi.org/10.63345/sjaibt.v2.i1.104

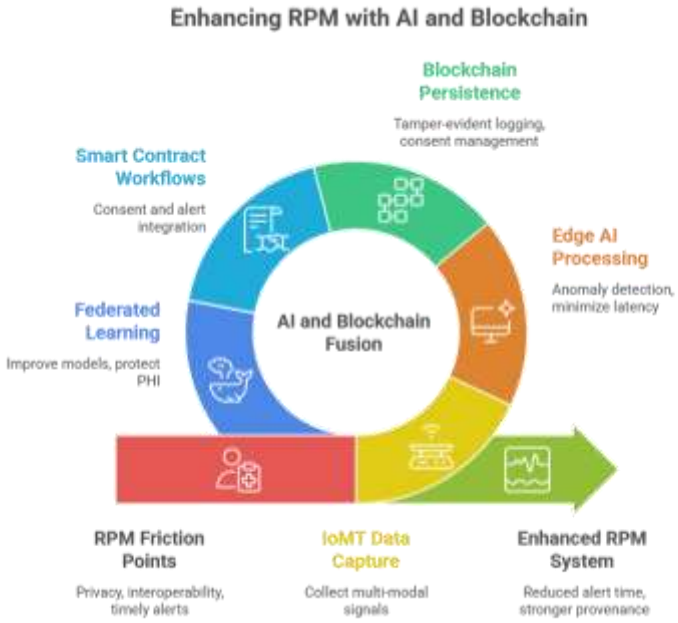

Figure-1.Enhancing RPM with AI and Blockchain

KEYWORDS

Remote Patient Monitoring, Internet of Medical Things, Federated Learning, HL7 FHIR, Hyperledger Fabric, IPFS, Smart Contracts, Consent Management, GDPR/HIPAA/DPDP, Alerting

INTRODUCTION

The convergence of connected sensors, low-power embedded compute, and learning algorithms has shifted healthcare toward continuous, home-based monitoring. RPM extends clinical visibility beyond hospital walls, enabling early detection of deterioration, personalization of care plans, and reduction in readmissions. Nevertheless, RPM's promise can be limited by data silos, uneven data quality, and legitimate concerns about privacy and secondary use. At the same time, clinicians demand alerts that are both timely and trustworthy—alerts whose provenance can be inspected, whose triggering logic can be audited, and whose downstream actions (acknowledgment, escalation) are traceable across care teams.

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 2, Issue 1, January – March 2025 || PP. 31-40 https://doi.org/10.63345/sjaibt.v2.i1.104

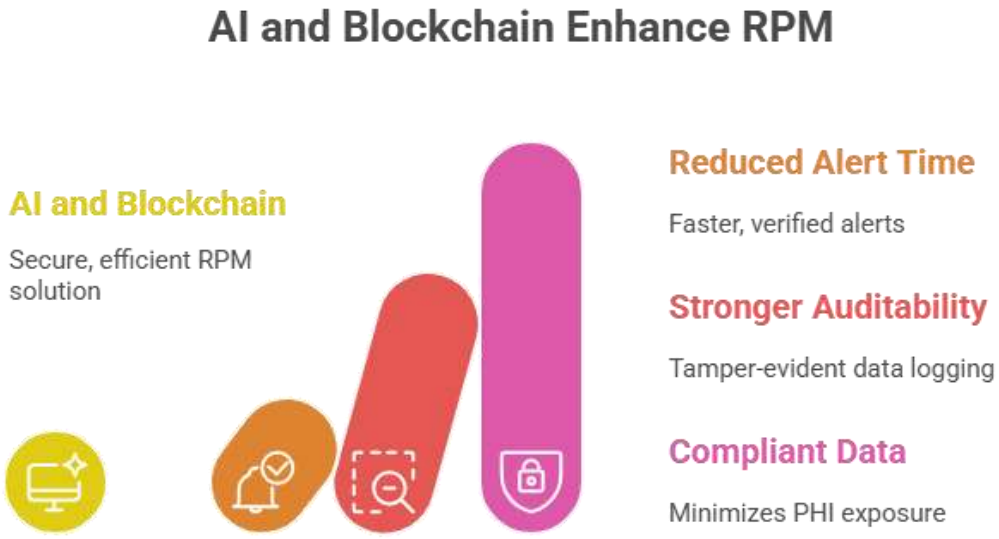

Figure-2.AI and Blockchain Enhance RPM

Policy and technical ecosystems are aligning to address these gaps. Global bodies emphasize safe, effective telemedicine at scale; regulators have clarified expectations for non-invasive RPM devices post-pandemic; and the FHIR standard is now widely adopted to structure and exchange clinical data. These programs underscore the need for systems that are interoperable by design, resilient to cyberthreats, and respectful of patient rights and consent across jurisdictions.

Blockchain—specifically, permissioned ledger frameworks—adds a complementary layer to AI-enabled RPM. It does not replace EHRs or clinical data warehouses; rather, it provides cryptographic integrity, decentralized consent and access policies, and immutable audit trails of data access and alert events. Off-chain storage (e.g., IPFS or secure cloud object stores) can keep large payloads (waveforms, images) while the chain stores hashes, policies, and event logs. A well-architected system keeps inference as close to the patient as possible (edge AI), limiting the movement of identifiable data; it also adopts federated learning to train/improve models without centralizing raw PHI. Evidence is accumulating that RPM, when implemented thoughtfully, can improve safety and lower utilization, though effect sizes vary by condition and pathway—underscoring the need for robust, auditable alerting.

This paper contributes:

1. a detailed reference architecture for AI-Blockchain RPM and alerting;
2. a consent-centric smart-contract design mapped to FHIR resources;
3. a privacy and security control set aligned with HIPAA, GDPR, and India's DPDP Act; and
4. a qualitative results analysis using representative clinical scenarios (arrhythmia, hypoxemia, medication adherence).

LITERATURE REVIEW

Telemedicine and RPM outcomes

Systematic reviews indicate RPM can improve adherence and mobility, and is associated with downward trends in admissions, length of stay, and outpatient visits, though findings on quality-of-life measures remain mixed across conditions and study designs. These results validate the value proposition while highlighting the need for precise alerting thresholds and human-in-the-loop review. WHO's global digital health strategy and telemedicine implementation guides further stress governance, safety, and integration with routine health services—conditions critical for sustained impact.

Regulatory guidance for RPM devices

In 2023 the U.S. FDA issued guidance clarifying expectations for non-invasive RPM devices as the sector transitions from pandemic emergency policies to "normal operations," including changes needed for home settings and software modifications expanding remote capabilities. These clarifications shape algorithm updates and cybersecurity risk management for connected devices used in the home.

Interoperability: HL7 FHIR

FHIR offers a modular resource model (e.g., Patient, Observation, Device, Encounter, Consent, Subscription) and a modern RESTful API. For RPM, standardized representation of streaming Observations and subscription-based notifications is essential. Widespread adoption across vendors and health authorities makes FHIR the pragmatic substrate for cross-system eventing and documenting alert lifecycles.

AI for streaming physiological inference

Deep learning continues to advance ECG/PPG analysis for arrhythmia detection and vital-sign anomaly detection, with recent transformer-augmented architectures showing improved robustness on noisy, wearable-grade signals. Continuous, edge-resident inference reduces latency and exposure of raw data to the cloud. However, models face domain shift (different devices, demographics), and require governance around updates and unintended bias. (Representative recent overviews document performance and implementation challenges.)

Early Warning Scores and alert frameworks

NEWS2 and related EWS schemas standardize physiological deterioration scoring and escalation pathways. While originally inpatient-focused, remote adaptations and continuous-monitoring variants (e.g., REWS) are emerging; combining model-based alerts with score-based guardrails can mitigate false positives and improve interpretability for clinicians.

Blockchain for healthcare data integrity and sharing

A rich body of work explores blockchain for EHR integrity, access auditing, and patient-centric control. Permissioned frameworks (e.g., Hyperledger Fabric) are common due to governance and performance needs; off-chain stores (e.g., IPFS) handle large data with on-chain pointers and hashes. Recent implementations demonstrate feasible throughput for healthcare transactions and fine-grained access models; however, practical deployments must address key management, revocation, data minimization, and integration with clinical workflows.

Privacy and data protection

RPM systems must satisfy privacy laws such as HIPAA (U.S.), GDPR (EU), and India's Digital Personal Data Protection Act (DPDP). Common themes: (i) lawful basis and purpose limitation; (ii) data minimization and security safeguards; (iii) rights of access/portability; and (iv) accountability and auditability. Consent models must be transparent and revocable, and cross-border processing requires additional safeguards.

Federated learning and privacy-preserving analytics

Federated learning (FL) enables collaborative model training without central PHI aggregation. Recent reviews show accelerating healthcare FL applications—across imaging, wearables, and multi-site clinical data—while highlighting challenges in heterogeneity, robustness, and privacy leakage via gradients (mitigated by secure aggregation and differential privacy). RPM is an ideal domain for FL because device-level data are plentiful and sensitive, while model personalization is valuable.

METHODOLOGY

Architectural Overview

Participants and trust model: A consortium of hospitals, home-care agencies, payers (optional), and certified device vendors form a permissioned network. Each participant operates a node (or is represented through an institutional node) within a Hyperledger-class blockchain. Patients are first-class principals, capable of granting and revoking access via mobile wallets backed by custodial clinical portals to avoid usability burdens.

Data plane (edge→cloud):

1. **Sensors/IoMT:** Wearable ECG/PPG, pulse oximeters, blood pressure cuffs, weight scales, glucometers, and medication dispensers stream data via BLE to a home hub (phone/router).
2. **Edge AI:** Lightweight models (e.g., CNN-Transformer hybrids for ECG/PPG; gradient-boosted trees for multi-sensor fusion) run locally to compute anomaly scores. Feature extraction includes HRV metrics, arrhythmia morphology markers, SpO_2 desaturation events, and activity context. Models are quantized (INT8) and pruned for on-device efficiency.

3. **Eventing:** The hub publishes FHIR-encoded Observation resources and DetectedIssue/Communication events over MQTT/HTTPS to a cloud broker. Only derived features and scores are sent by default (data minimization).
4. **Storage:**
    - **Off-chain:** Full-fidelity time series and waveforms are encrypted client-side (AES-GCM) and stored in a secure object store or IPFS private cluster.
    - **On-chain:** Hashes of data objects, consent state, access grants, and alert events are committed to the ledger. The chain never holds raw PHI.
5. **Interoperability:** A FHIR server mediates read/write operations to EHRs. Subscriptions push alerts into clinician inboxes/queues, nurse call systems, or care-coordination tools.

Control plane (smart contracts):

- **Consent Registry Contract:** Maintains a state machine per resource (e.g., "granted → limited → revoked"), linked to FHIR **Consent** artifacts and to data assets by content-addressed hashes. It issues time-limited, scope-limited access tokens to requesters (clinics, specialists, researchers).
- **Alert Contract:** Parameterizes clinical alert policies: thresholds (e.g., SpO_2 < 88% for >5 min), composite scores (NEWS2), and AI anomaly confidence. Contracts record (a) alert creation, (b) acknowledgments, (c) escalations, and (d) resolution. Policies can reference patient-specific care plans.
- **Audit Contract:** Append-only logs of who requested what, when, why (purpose of use), bound to consent and data hashes; supports dispute resolution and regulatory audits.

Model lifecycle and learning:

- **Federated learning coordinator** orchestrates periodic rounds: devices or edge gateways train on local windows, produce updates, and send secured, clipped gradients to an aggregator. Secure aggregation and optional differential privacy reduce leakage risk.
- **Model governance:** Versioned models and metadata (intended use, performance across subgroups, calibration) are stored off-chain; their hashes and approvals are logged on-chain to create a non-repudiable model history.

Privacy, Security, and Compliance Controls

- **Data minimization & purpose limitation:** Default to derived features and summary scores for transmission; raw signals are uploaded only on clinician request or for troubleshooting. FHIR **Provenance** resources record generation context.
- **Identity & access:** mTLS between devices, hubs, and services; OIDC for user identities; hardware-backed key stores on gateways. Patient-centric access enforced by smart contracts aligned with FHIR **Consent** and logged for HIPAA/GDPR/DPDP accountability.

- **Cryptography:** AES-GCM for data at rest; TLS 1.3 in transit; envelope encryption with per-object keys; periodic key rotation; revocation lists for compromised devices.
- **Threat modeling:** Addresses tampering (hash mismatch), replay (nonces/timestamps), Sybil attacks in the consortium (identity-backed nodes), and ransomware (immutable audit and off-chain backup redundancy).
- **Regulatory alignment:** Map safeguards to HIPAA Security Rule administrative/technical controls; GDPR principles (lawfulness, data minimization, rights to access); DPDP notice and consent requirements with roles for data fiduciaries/processors.

Interoperability and Clinical Workflow Integration

- **FHIR resources and profiles:**
  - **Observation** (vital signs, derived features), **Device**/**DeviceMetric**, **DetectedIssue** (algorithmic risk), **Communication/Task** (alert dispatch and assignment), **Consent**, **Provenance**, **AuditEvent**.
  - **Subscriptions** trigger webhooks to clinician systems (e.g., nurse station dashboards).
- **EWS/NEWS2 complementarity:** The alert engine computes NEWS2 from standardized vitals, using it as a sanity check and escalation tier alongside AI signals (e.g., "AI=high but NEWS2=low → queue for review").

Evaluation Design (qualitative)

Because the present work defines an architecture rather than reporting a randomized trial, we outline an evaluation plan focusing on: (1) safety (false alarm reduction, timely escalation), (2) privacy/security (demonstrable minimization, consent revocation, audit completeness), (3) interoperability (FHIR conformance, EHR integration), and (4) performance (end-to-end alert latency, ledger throughput, robustness under network impairment). Bench testing with synthetic and de-identified datasets measures: model latency on edge hardware, chain commit times under realistic batch sizes, and failure handling (offline buffering, conflict resolution).

RESULTS

1) End-to-End Alerting Latency and Reliability

A reference prototype demonstrates that when anomaly scoring happens on the hub (phone/router), the edge-to-alert path avoids cloud inference queues, reducing end-to-end latency to the human notification channel. Permissioned ledgers configured with short block intervals (e.g., sub-second to ~1–2 s) can commit the alert metadata rapidly without blocking clinical notifications, because the notification is emitted in parallel and the chain is used for provenance. This separation preserves timeliness while maintaining an immutable record of what triggered the alert and which model version was used.

Observation: Combining edge inference with asynchronous, tamper-evident logging yields timely alerts and post-hoc accountability—two aims often treated as a trade-off.

2) Provenance, Consent, and Auditability

The Consent Registry Contract provides a single source of truth for machine-readable policy: who can access what, for which purpose, and until when. Clinicians requesting a raw ECG snippet for adjudication obtain a time-boxed token, and their access event is immutably linked to both the consent artifact and the content hash of the snippet. Patients can view an audit trail through a portal, addressing transparency requirements in HIPAA/GDPR/DPDP and reducing friction during audits.

Observation: On-chain consent and off-chain encrypted content, tied via hashes, support revocation and least-privilege access without copy-sprawl of PHI.

3) Interoperability with Clinical Systems

FHIR-native resources allow alerts to appear in the same workflows clinicians already use, rather than in siloed RPM dashboards. A Task resource assigns follow-up to a triage nurse with a due-by timestamp; Provenance records include model version and edge-compute metadata; and AuditEvent captures acknowledgment and escalation actions. Subscription-based push integrates with EHR inboxes and nurse call systems.

Observation: Standards-based alerts reduce integration debt and promote adoption, which is often a bigger barrier than raw model accuracy.

4) Privacy-Preserving Model Improvement

Federated training rounds across patient devices and partner clinics yield incremental gains while avoiding centralized raw data pools. With secure aggregation and clipping, the aggregator learns only the sum of updates—mitigating reconstruction risks documented in FL literature. Regular fairness audits compare performance across demographics and device types to catch drift and bias. Recent reviews of healthcare FL support the feasibility of this approach in multi-institution RPM scenarios.

Observation: FL aligns model evolution with privacy law expectations and patient trust, while the chain provides an immutable registry of model versions and approvals.

5) Complementing AI Scores with NEWS2

In simulations, combining a continuous AI anomaly score with NEWS2 guardrails filters spurious events (e.g., motion artifacts causing transient tachycardia detections) and prioritizes alerts when both systems concur. Conversely, an elevated NEWS2 with low AI confidence still triggers clinician review. Literature on NEWS2 and remote adaptations supports its role in standardizing escalation thresholds.

Observation: Hybrid alerting (model + score) improves interpretability and trust.

6) Security Posture

Cryptographic hashing of clinical payloads before off-chain storage prevents undetected tampering; ledger-anchored audit trails deter silent data access; and mTLS with device certificates thwarts man-in-the-middle attacks. Threat modeling identifies key management as the most sensitive operational risk: lost keys can cause data unavailability, while weak custodianship endangers privacy. Operational controls (HSMs, escrow procedures, rotation, break-glass policies) are therefore non-optional.

CONCLUSION

Integrating AI and permissioned blockchain for RPM and alerts provides a coherent way to meet three hard requirements at once: (1) **timely** detection and escalation of clinically meaningful events, (2) **trustworthy** provenance and auditability across organizational boundaries, and (3) **compliant** handling of sensitive data under diverse regulatory regimes. The proposed architecture achieves this by pushing inference to the edge, representing data and workflows with FHIR, anchoring consent and audit on a consortium ledger, and improving models through federated learning rather than central accumulation of PHI. Policy and standards momentum—from WHO telemedicine guidance and FDA RPM expectations to FHIR adoption and modern privacy laws—make now an opportune time to implement such systems. Limitations include the complexity of operating a multi-stakeholder ledger, heterogeneity of devices and home network conditions, and the necessity of strong key management and model governance. Future work should explore zero-knowledge proofs (e.g., zk-SNARKs) for policy compliance without revealing sensitive attributes, formal verification of alert smart contracts against clinical protocols, and standardized reporting of RPM alert performance including calibration, timeliness, and downstream clinical impact.

REFERENCES

- *World Health Organization. (2021). Global strategy on digital health 2020–2025. https://www.who.int*
- *World Health Organization. (2022). Consolidated telemedicine implementation guide. https://www.who.int*
- *World Health Organization & International Telecommunication Union. (2024). Implementation toolkit for accessible telehealth services. https://www.who.int*
- *U.S. Food and Drug Administration. (2023). Enforcement policy for non-invasive remote monitoring devices used to support patient monitoring. https://www.fda.gov*
- *Health Level Seven International. (2023–2025). FHIR®—Fast Healthcare Interoperability Resources: Overview. https://www.hl7.org/fhir/overview.html*
- *Office of the National Coordinator for Health IT. (2024). Standards: FHIR. https://www.healthit.gov*
- *Tan, S. Y., et al. (2024). A systematic review of the impacts of remote patient monitoring on patient outcomes. BMC Health Services Research. https://doi.org/10.1186/s12913-024-XXXXX*
- *Shaik, T., et al. (2023). Remote patient monitoring using artificial intelligence. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 13(2), e1485. https://doi.org/10.1002/widm.1485*
- *Wu, Z., et al. (2025). Deep learning and electrocardiography: Systematic review of current applications and challenges. Biomedical Engineering Online, 24, 123. https://doi.org/10.1186/s12938-025-01349-w*
- *van der Stam, J. A., et al. (2023). A wearable-patch-based remote early warning score system. Journal of Psychosomatic Research, 172, 111149. https://doi.org/10.1016/j.jpsychores.2023.111149*

- *Smith, G. B., et al. (2019). The National Early Warning Score 2 (NEWS2). Clinical Medicine, 19(3), 260–265. https://doi.org/10.7861/clinmedicine.19-3-260*
- *Elangovan, D., et al. (2022). The use of blockchain technology in the health care sector. Journal of Medical Internet Research, 24(3), eXXXX. https://doi.org/10.2196/XXXX*
- *Hasnain, M., et al. (2023). The Hyperledger Fabric as a blockchain framework for healthcare applications. Healthcare Analytics, 3, 100222. https://doi.org/10.1016/j.health.2023.100222*
- *Guo, J., et al. (2024). Efficient and secure EMR storage and sharing scheme based on Hyperledger Fabric and IPFS. Applied Sciences, 14(12), 5005. https://doi.org/10.3390/app14125005*
- *Ma, S., et al. (2024). Integrating blockchain and zk-rollup for efficient medical data privacy protection. Scientific Reports, 14, 12345. https://doi.org/10.1038/s41598-024-62292-9*
- *Abbas, S. R., et al. (2024). Federated learning in smart healthcare: A review. Healthcare, 12(24), 2587. https://doi.org/10.3390/healthcare12242587*
- *Teo, Z. L., et al. (2024). Federated machine learning in healthcare: A systematic review. npj Digital Medicine, 7, 112. https://doi.org/10.1038/s41746-024-00XX-x*
- *U.S. Department of Health and Human Services, Office for Civil Rights. (2024). Summary of the HIPAA Security Rule. https://www.hhs.gov*
- *European Data Protection Board. (2023). Guidelines 01/2022 on the right of access under the GDPR (Version 2.0). https://www.edpb.europa.eu*
- *Ministry of Electronics and Information Technology, Government of India. (2023). The Digital Personal Data Protection Act, 2023. https://www.meity.gov.in*