



# Architectural Patterns for Building Scalable Cloud Platforms for Continuous Remote Cardiac Monitoring of Patients

Oleksandr Kydiuk

Lead Software Engineer, GlobalLogic, Kyiv, Ukraine.

## Abstract

*The article examines a set of architectural patterns that enable the construction of scalable cloud platforms for continuous remote cardiac monitoring of patients amid growing demand for timely detection of cardiovascular disorders. The relevance of the study is determined by the limitations of traditional outpatient monitoring methods, which often record clinically significant arrhythmia episodes with substantial delay, thereby reducing the effectiveness of early intervention and increasing the risk of adverse outcomes. The aim of the work is to develop an architectural model of a cloud platform capable of providing reliable, secure, and near-real-time processing of streaming ECG data within the Internet of Medical Things environment. The scientific novelty of the article lies in substantiating the concept of seamless monitoring by synthesizing edge-based signal filtering, Bluetooth gateway data transmission, and a microservice-based, event-driven cloud infrastructure. As key conclusions, it is shown that the use of a Bluetooth gateway makes it possible to eliminate the critical gap between ECG recording and interpretation, that a microservice architecture in combination with Kafka and Raft ensures horizontal scalability and fault tolerance, and that a hybrid edge-cloud processing model improves analytical accuracy and reduces the latency of processing critical events. The article will be useful for researchers, architects of medical IT systems, engineers of IoMT platforms, and specialists in digital healthcare.*

**Keywords:** Remote Cardiac Monitoring, Cloud Architecture, Internet of Medical Things, Scalability, Bluetooth Gateways, Microservices, Event-Driven Processing, Machine Learning.

## INTRODUCTION

Modern healthcare is undergoing a digital transformation driven by the need of efficiently managing chronic diseases, among which cardiovascular diseases are the leading cause of death and disability (Freihat et al., 2025). The standard practice for cardiovascular disease diagnosis is based on episodic testing in a clinical environment upon prescription, or short-term Holter monitoring (Carrington et al., 2022). It is less effective in detecting intermittent arrhythmias such as paroxysmal atrial fibrillation, which may occur infrequently but which are still associated with a high risk of thromboembolism.

The drawback is the time lapse between the event and the interpretation. In the classical monitoring system, data are stored in the device for 7 to 14 days. In order to receive an analysis, the patient has to return to the medical center to upload and analyze the data (Himmelreich et al., 2022). Because the delay can take several weeks, diagnosis and treatment of any underlying condition may be impacted. Thus,

a need arises for the development of architectural solutions capable of ensuring continuous automated transmission and processing of high-frequency cardiac signals in a mode that approximates real time as closely as possible.

The aim of the study is to develop an architectural model for a scalable cloud platform for continuous remote cardiac monitoring, ensuring prompt, reliable, and secure processing of electrocardiograms in near-real time.

To achieve the aim of the study, the following objectives are addressed:

to consider the limitations of traditional methods of outpatient cardiac monitoring;

to examine contemporary approaches to building Internet of Medical Things platforms for ECG data processing;

to determine architectural solutions that ensure scalability, fault tolerance, and low latency;

to characterize the role of microservice, event-driven, and hybrid edge-cloud architecture;

**Citation:** Oleksandr Kydiuk, "Architectural Patterns for Building Scalable Cloud Platforms for Continuous Remote Cardiac Monitoring of Patients", Universal Library of Engineering Technology, 2026; 3(1): 104-108. DOI: <https://doi.org/10.70315/uloap.ulete.2026.0301018>.

to investigate methods of signal filtering, protection of medical data, and intelligent event processing.

The object of the study is remote cardiac monitoring systems for patients implemented using cloud technologies and the Internet of Medical Things infrastructure. The subject of the study is architectural patterns, computational organization models, and mechanisms for processing, transmission, and protection of ECG data that ensure the effective functioning of scalable cloud platforms for remote cardiac monitoring.

The scientific novelty of the work lies in the synthesis of the concept of seamless monitoring, which integrates edge computing methods for primary data filtering and a cloud microservice infrastructure for in-depth analysis, enabling a latency of less than 250 ms in processing critical events.

**MATERIALS AND METHODOLOGY**

The methodological foundation of this study is constructed through a combination of industrial case-study data analysis and qualitative synthesis of the theoretical foundations of cloud computing. As primary material, data from a remote cardiac monitoring platform development project were used; over ten years of operation, this platform accumulated tens of billions of hours of ECG recordings and serves hundreds of thousands of patients annually.

To ensure scientific depth, a literature review was conducted for the period 2020–2025. The search was conducted across the repositories Scopus, Web of Science, IEEE Xplore, and ACM Digital Library using the keywords: IoMT cloud architecture, ECG streaming scalability, Bluetooth gateway healthcare. The final analysis included sources describing event-driven architectures for medical data processing, distributed consensus algorithms in the context of IoMT, methods of integrating deep learning into cloud pipelines, and comparative characteristics of data transmission protocols.

**Table 1.** Decomposition of the architecture levels of the cardiac monitoring platform

Layer	Technology Stack	Primary Task	Applied Pattern
Edge Layer	STM32, BLE 5.0, TinyML	Raw data collection, artifact filtering, local buffering	Edge Intelligence
Ingestion Layer	API Gateway, Apache Kafka	Device authentication, message queue management	Event-Driven Ingestion
Processing Layer	Apache Flink, Spark Streaming	Real-time arrhythmia detection, feature extraction	Microservices Architecture
Storage Layer	NoSQL (MongoDB), S3, PostgreSQL	Distributed storage of time-series data and reports	Data Partitioning & Sharding
Insight Layer	WebSocket, RESTful API	Data visualization, interactive dashboards for physicians	Publish-Subscribe

The application of microservice architecture at the processing level enables isolating resource-intensive tasks, such as deep ECG analysis by neural networks, from data ingestion services. This ensures independent scaling: when the number of patients increases, the system automatically increases the number of instances of the detection service without affecting the stability of the data ingestion gateway. Figure 1 below is the developed conceptual scheme of data flows in the system.

The analytical approach included a comparative analysis of various network topologies and an assessment of their impact on the energy consumption of wearable devices and the reliability of data delivery. The mathematical substantiation of system reliability was based on median filtering models under sensor redundancy.

**RESULTS AND DISCUSSION**

One of the central results of the study is the demonstration of the advantage of an architecture that uses an intermediate Bluetooth gateway to bridge the wearable sensor and the cloud platform. Unlike systems that require the physical return of the device, this model provides daily automated data uploads without patient involvement.

The technical implementation of this pattern uses Bluetooth Low Energy for sensor-gateway communication. This tradeoff can be influenced by strict power consumption requirements, as in wearable ECG patches lasting 7 to 14 days, and wireless protocols like Bluetooth Low Energy are often selected for their smaller size and low-power requirements over alternatives such as cellular. This is particularly important for higher frequency transmissions (Schuhmacher et al., 2023). Once the patient is back within range of the home router, the device creates a session and uploads the data. The data is encrypted before transmission to the cloud via a reliable Wi-Fi or wired internet connection.

This model reduces the delay time of clinical interpretation. If, in the traditional scheme, the physician saw the data only after 14–21 days, then in the gateway-based architecture, access to current information is provided within several hours, and critical alerts can be generated in near-real time.

To ensure resilience under peak loads, when thousands of devices simultaneously initiate data upload, for example, during morning hours, the platform must be organized hierarchically. Table 1 shows the decomposition of the architectural levels of the cardiac monitoring platform.

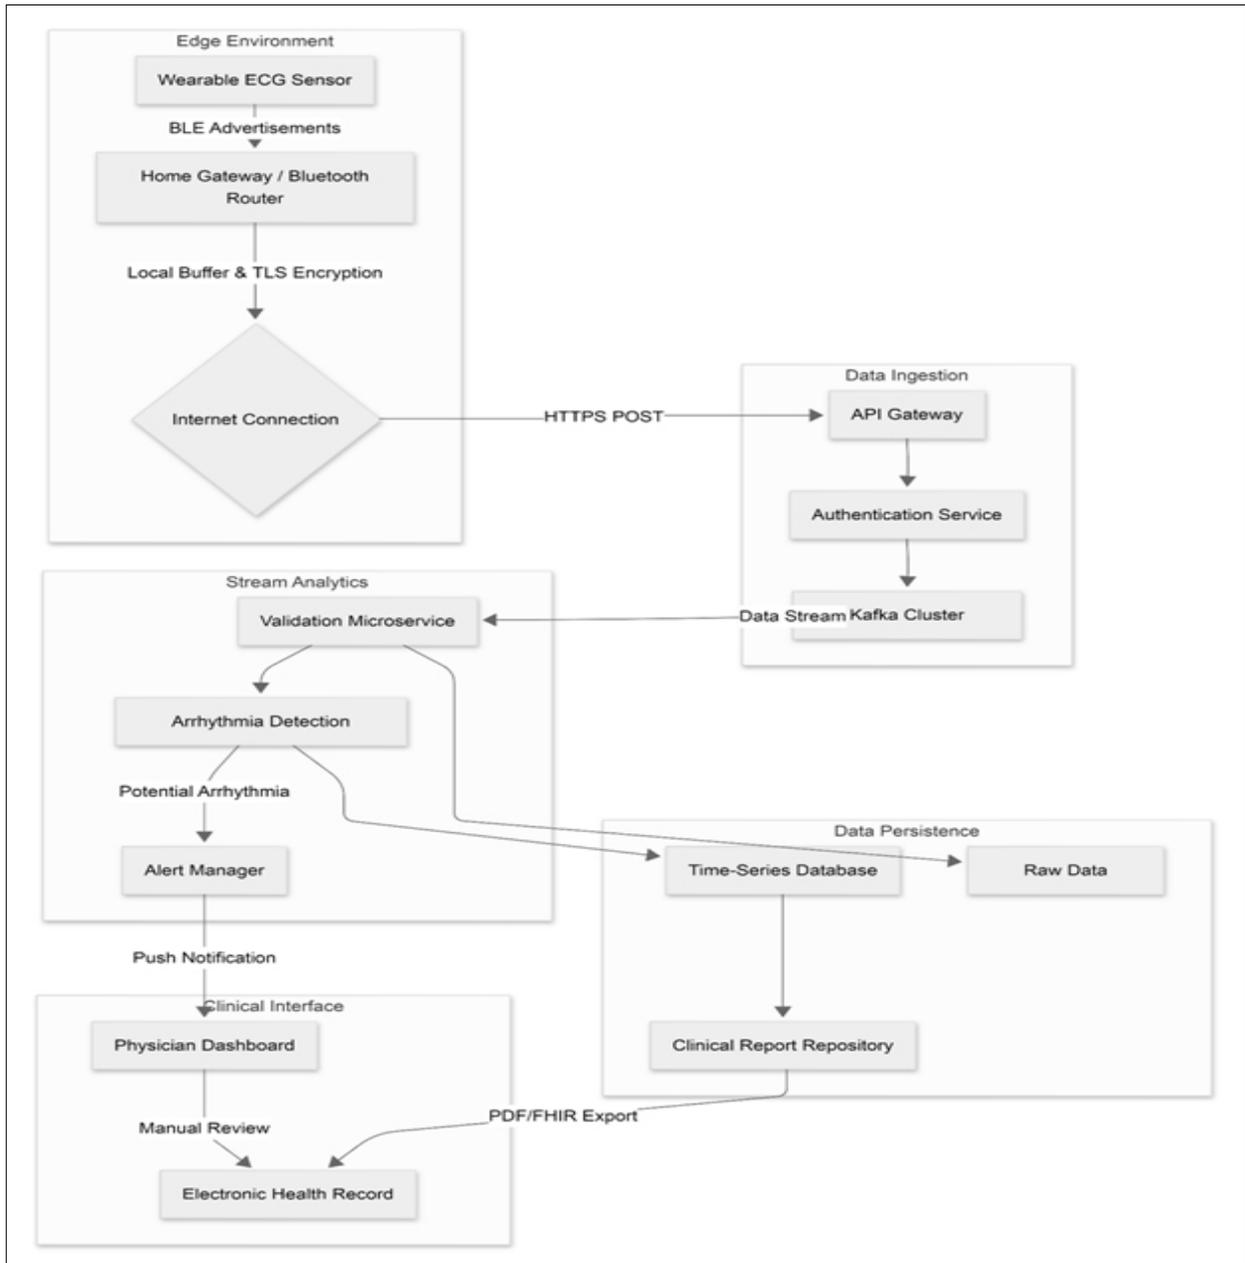


Fig. 1. Architecture of a distributed cloud platform for continuous monitoring

The cloud platform succeeded mostly due to its ability to maintain a high throughput at a low latency. Experimental tests show that Kafka (Arafat, 2025) is able to achieve millions of messages per second throughput at tens of milliseconds latency, making it a reliable solution.

The IoMT is a distributed system with massive amounts of data generated from the sensors, so the consensus algorithm is critical to ensuring the integrity of the data (Hao et al., 2025). In the architecture, the Raft protocol was used, which has been shown to achieve 1.26 million transactions per second with up to 1000 nodes and a 92.3% reduction in consensus latency. This ensures that the ECG data packets from the wearable device are properly stored and processed, even in the event of a server failure.

A hybrid signal-processing model combining edge and cloud

filtering is substantiated. ECG signals are highly sensitive to various types of noise and artifacts, especially those caused by patient movement. In wearable cardiac monitoring systems, such artifacts are a principal source of signal distortion and can substantially degrade data quality and diagnostic accuracy (An et al., 2022).

At the edge device level, it is expedient to use lightweight algorithms, such as median filtering with a 3-of-5 scheme, to eliminate obvious interference before data transmission to the cloud, thereby conserving channel bandwidth. In the cloud, heavyweight deep learning models are used, such as 1D-CNNs for QRS complex detection and LSTMs for analyzing rhythmic patterns over time. The accuracy of such systems in four-class arrhythmia classification reaches 96.5% (Sanjeev et al., 2025). The pattern of intelligent processing of ECG signals is illustrated in Figure 2.

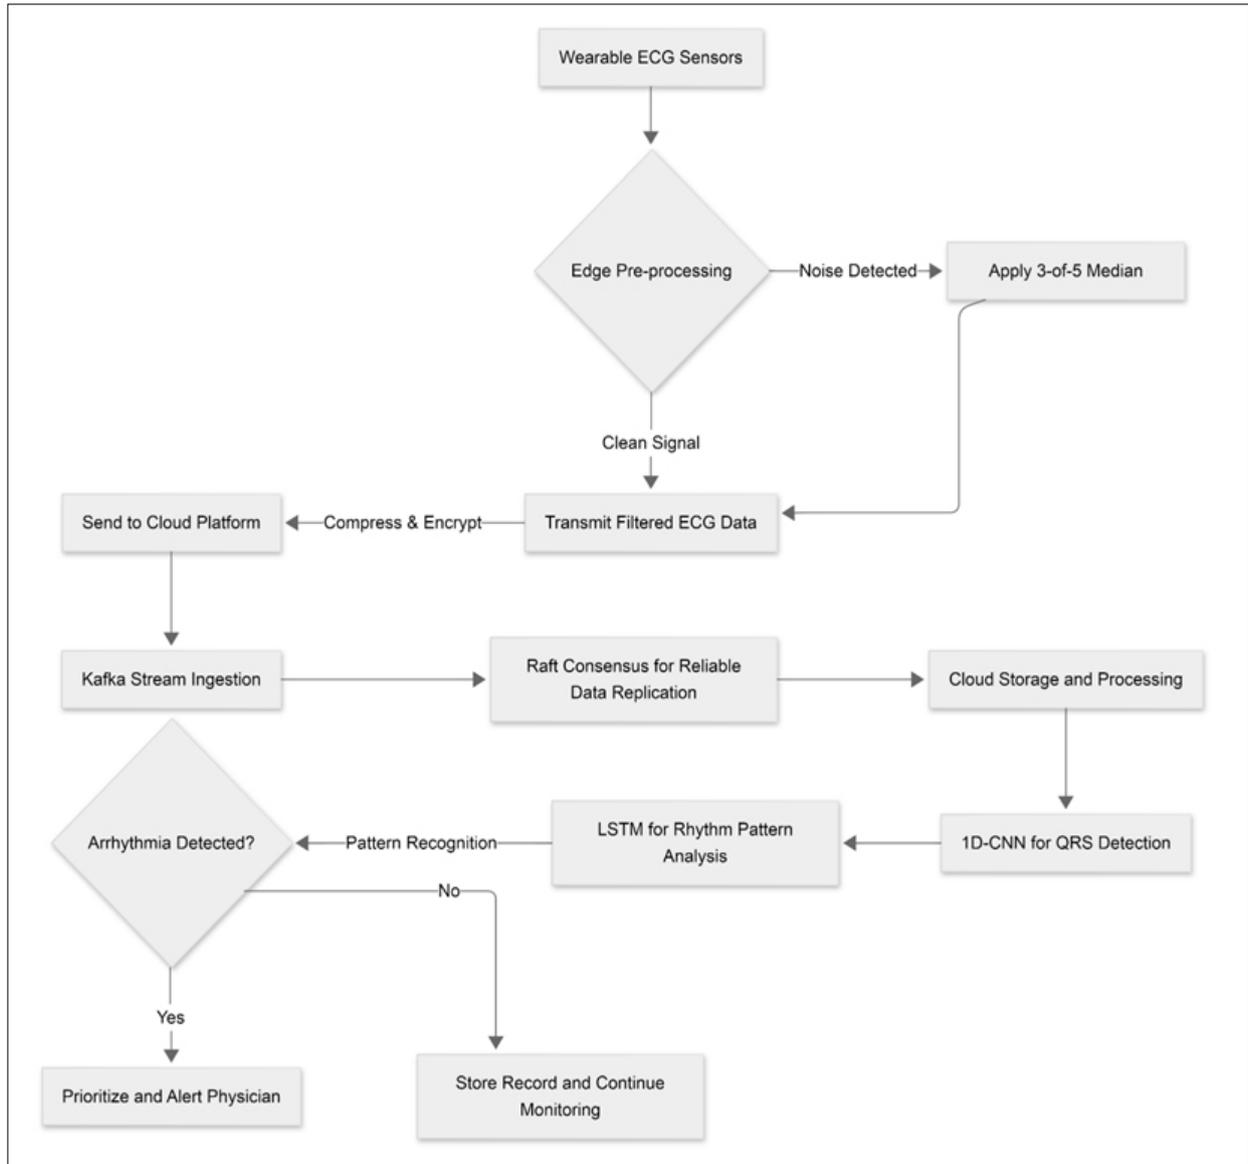


Fig. 2. Pattern of intelligent processing of ECG signals

The choice of intermediary for data transmission directly affects compliance and patient consent to technology use. At the same time, research shows that older people and people with lower technology literacy prefer a dedicated Bluetooth router to a smartphone (Wu & Lim, 2024).

As a case study, the Bluetooth router model allowed the company to do national expansion and reach people who had difficulty using digital data collection because of smartphone setup complexity.

The security of IoMT cloud platforms is a fundamental requirement enshrined in HIPAA and GDPR standards. The following patterns must be implemented in the architecture. End-to-End Encryption presupposes that data are encrypted at the sensor level and decrypted only within a protected cloud enclave. Such an approach precludes access to Personal Health Information at intermediate nodes, including the gateway and smartphone.

Post-quantum resilience involves using algorithms such

as Kyber to protect encryption keys in the long term. This is necessary to preserve the cryptographic security of the system as next-generation computational capabilities develop.

Storage fault tolerance is ensured through triple data replication and geographic redundancy. Such a storage organization is intended to prevent data loss in the event of catastrophic failures in the data center.

The principal risk remains alert fatigue, a situation in which medical personnel cease to respond adequately to notifications due to their excessiveness. In response, a smart triage layer is added to the architecture to sort the events in terms of clinical relevance and algorithm confidence in detection.

The objective of this study is the architectural level related to continuous cardiac monitoring and does mainly serve the purpose to touch upon the outer frame of a scalable building block with respect to reliability, responsiveness,

and data protection technology. Clinical use cases and regulatory environments of different countries as well as different living environments, which are liable to impact both the robustness of signal transmission and the quality of user experience, are not covered here. Furthermore, the proposed model draws on a synthesis of case data, literature, and engineering practices, ensuring that its implementation potential is fully realized through subsequent adaptation to the specific device configuration, patient sample profile, and medical information system framework. This gives the study's results the character of a universal architectural foundation, suitable for further customization depending on clinical, organizational, and infrastructural parameters.

### CONCLUSION

The study of architectural patterns for scalable cloud-based cardiac monitoring systems yields several conclusions of both theoretical and practical significance.

First, an architecture employing a Bluetooth gateway is optimal for long-term continuous monitoring, as it effectively resolves the tension between minimizing the energy consumption of the wearable device and the need for prompt data transmission to the cloud. Second, the transition to an event-driven microservice infrastructure based on Apache Kafka and the Raft consensus algorithm ensures the necessary level of horizontal scalability, enabling the system to process streams from hundreds of thousands of devices with sub-second latency.

The practical significance of the work is confirmed by the results of implementation in the case study, where it was possible to reduce the physician's access time to clinical data from weeks to several hours, thereby directly affecting patient safety. The original authorial contribution, in the form of a hybrid signal-processing model, enables effective noise and motion artifact reduction, thereby increasing the diagnostic value of monitoring.

Future research in this area should focus on ensuring semantic interoperability and on integrating cardiac monitoring data with electronic medical records to create comprehensive predictive models of patient health status.

### REFERENCES

1. An, X., Liu, Y., Zhao, Y., Lu, S., Stylios, G. K., & Liu, Q. (2022). Adaptive Motion Artifact Reduction in Wearable ECG Measurements Using Impedance Pneumography Signal. *Sensors*, 22(15), 5493. <https://doi.org/10.3390/s22155493>
2. Arafat, J. (2025). Next-Generation Event-Driven Architectures: Performance, Scalability, and Intelligent Orchestration Across Messaging Frameworks. *ArXiv*. <https://doi.org/10.48550/arxiv.2510.04404>
3. Carrington, M., Providência, R., Chahal, C. A. A., Ricci, F., Epstein, A. E., Gallina, S., Fedorowski, A., Sutton, R., & Khanji, M. Y. (2022). Monitoring and Diagnosis of Intermittent arrhythmias: evidence-based Guidance and Role of Novel Monitoring Strategies. *European Heart Journal Open*, 2(6), oeac072. <https://doi.org/10.1093/ehjopen/oeac072>
4. Freihat, O., Sipos, D., Aamir, M., & Kovacs, A. (2025). Global burden and future projections of non-communicable diseases (2000–2050): Progress toward SDG 3.4 and disparities across regions and risk factors. *PLoS ONE*, 20(12), e0336036. <https://doi.org/10.1371/journal.pone.0336036>
5. Hao, L., Wang, R., Wang, X., Yue, X., Tariq, N., & Sajid, A. (2025). Post-quantum-inspired scalable blockchain architecture for internet hospital systems with lightweight privacy-preserving access control. *PLOS One*, 20(12), e0332887. <https://doi.org/10.1371/journal.pone.0332887>
6. Himmelreich, J. C. L., Lucassen, W. A. M., Coutinho, J. M., & Harskamp, R. E. (2022). 14-day Holter monitoring for atrial fibrillation after ischemic stroke: The yield of guideline-recommended monitoring duration. *European Stroke Journal*, 8(1), 157–167. <https://doi.org/10.1177/23969873221146027>
7. Sanjeev, N. S., Yashvardhan, A., & Kacham, S. P. (2025). Detection of Heart-Diseases Through Cloud-Enhanced ECG Monitoring. *Proceedings of the 3rd International Conference on Futuristic Technology*, 57–64. <https://doi.org/10.5220/0013586600004664>
8. Schuhmacher, L., Pollin, S., & Sallouha, H. (2023). ecoBLE: A Low-Computation Energy Consumption Prediction Framework for Bluetooth Low Energy. *ArXiv*. <https://doi.org/10.5555/3639940.3639941>
9. Wu, C., & Lim, G. G. (2024). Investigating older adults users' willingness to adopt wearable devices by integrating the technology acceptance model (UTAUT2) and the Technology Readiness Index theory. *Frontiers in Public Health*, 12. <https://doi.org/10.3389/fpubh.2024.1449594>