Intelligent Environments 2025: Combined Workshop Proceedings
R. Aquino Santos and S. Faquiri (Eds.)
© 2025 The Authors.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/AISE250009

Towards a Unified Architecture for Remote Patient Monitoring

Federico BERGAMINI, Giovanni Donato GALLO and Daniela MICUCCI
University of Milano - Bicocca, Department of Informatics, Systems and Communication, Milan, Italy

Abstract. The increasing adoption of digital healthcare technologies is reshaping patient care, enabling intelligent environments that support both healthcare professionals and patients. However, challenges such as data heterogeneity, interoperability, security, and regulatory compliance hinder seamless integration into existing healthcare systems. This paper proposes a modular and scalable architecture for remote patient monitoring, designed to standardize data exchange, ensure secure communication, and enhance adaptability. The system leverages HL7 FHIR to facilitate interoperability across diverse healthcare platforms while incorporating privacy-preserving mechanisms, including encryption, authentication, and compliance with GDPR and HIPAA regulations. To validate its effectiveness, the proposed architecture is applied to cardiac surgery monitoring, integrating diverse devices and enabling personalized reporting to support patient care decisions.

Keywords. Remote patient monitoring, IoMT, interoperability, HL7 FHIR, security and compliance

1. Introduction

Digital healthcare platforms assist physicians in diagnostics and support patients in managing therapies and rehabilitation. However, many healthcare systems operate with proprietary data formats and communication protocols, making it difficult to integrate them into a unified solution. The data collected is often heterogeneous, complicating standardization. The widespread use of fitness trackers and wearable devices offers new opportunities for remote patient monitoring, but the large volume of data generated can be overwhelming and not always useful for diagnosis or treatment. Furthermore, each device has its own data format, increasing the complexity of healthcare data management. Security, reliability, and availability are crucial when dealing with healthcare data. Platforms must comply with privacy regulations such as GDPR and HIPAA to ensure data protection. Additionally, healthcare platforms must adapt to patients' heterogeneity, with personalized treatment plans and the ability to self-adapt, especially in cases of low device battery, by either replacing it or adjusting energy consumption.

This paper proposes MIRACLE (Modular Interoperable Reliable Architecture for CLinical Efficiency), a unified architecture for digital healthcare platforms, focusing on the following key aspects:

- **Data Integration and Interoperability.** Collecting and standardizing heterogeneous, multimodal data to ensure seamless exchange across diverse systems.

- **Standardized Communication.** Adopting syntactic and semantic standards, along with standardized protocols and service interfaces to facilitate data sharing.
- **Security & Compliance.** Implementing robust security measures, such as encryption, authentication, and continuous system monitoring, while ensuring compliance with privacy regulations like GDPR.
- **Personalization & Self-Adaptability.** Designing the system to support various medical applications and adapt to individual patient needs, ensuring personalized experiences for both patients and healthcare providers.

A case study on cardiac surgery monitoring validates the proposed architecture.

2. Related Work

Recent research in digital healthcare has focused on key aspects such as interoperability, multimodal and heterogeneous data integration, security, and legislation compliance. However, most proposed solutions address only one or a few of these dimensions, with few offering a comprehensive approach that covers all aspects simultaneously.

Several studies address interoperability and data integration. For instance, Hornback et al. [1] propose an Extraction, Transform, Load (ETL) pipeline that harmonizes multi-site, multimodal data into FHIR standardized structures, ensuring end-to-end security and interoperability through the SMART-on-FHIR standard. Similarly, Marfoglia et al. [2] present an ETL pipeline with a templating conversion strategy via FHIR Mapping Language. Li et al. [3] explore the use of Large Language Models (LLMs) for converting free text from practitioners into structured FHIR data, enhancing interoperability and data harmonization.

In terms of interoperability and security, Rindal et al. [4] integrate medical information into an Electronic Health Record (EHR) system using the SMART-on-FHIR standard, employing OAuth2 and OpenID Connect for authentication. Dos Santos et al.[5] propose a middleware-based FHIR architecture for stroke care, which addresses security, communication, and integration issues. Alamri et al. [6] introduce a blockchain-based electronic health wallet using IoT and FHIR for interoperability, ensuring GDPR compliance through blockchain and privacy layers.

Regarding legislation compliance, healthcare systems must comply with regulations like GDPR in Europe and HIPAA in the U.S. Chatterjee et al. [7] design a health coaching system using FHIR and SNOMED-CT for semantic interoperability, ensuring GDPR compliance with Norwegian Services for Sensitive Data (TSD) for secure data management. Vasileiou et al. [8] present a FHIR-compliant infrastructure using an Entity-Relationship semantic model, incorporating role-based access control, data encryption, and pseudonymization to comply with GDPR principles. Raso et al. [9] propose an open-source solution for de-identification of FHIR data, supporting anonymization and pseudonymization to meet regulatory requirements.

While these solutions contribute to the advancement of digital healthcare, they typically focus on specific challenges, often addressing interoperability and security, but not always integrating data processing pipelines for multimodal and heterogeneous data. Our approach aims to bridge these gaps by proposing a comprehensive digital healthcare platform that integrates multimodal data sources, ensures both semantic and syntactic interoperability, provides robust security, and complies with GDPR.

3. Key Features and Requirements

This section outlines the features and requirements that shaped the architecture design.
Interoperability. Interoperability in healthcare platforms is crucial for data exchange and collaboration across different systems. It involves three levels: *syntactic interoperability*, ensuring standardized data formats and communication protocols; *semantic interoperability*, achieved through controlled vocabularies and ontologies for precise data interpretation; and *organizational interoperability*, which aligns processes and policies for effective system collaboration. The HL7 FHIR standard addresses all three levels, ensuring data consistency, controlled coding systems (e.g., SNOMED-CT) to uniquely identify healthcare resources, and secure, standardized workflows.
Multimodality and heterogeneity. A key challenge is the management of multimodal and heterogeneous data originating from various sources. Consequently, data harmonization is essential to ensure that medical data, spanning structured electronic health records (EHRs) to sensor-derived information, can be effectively standardized and integrated. The flexibility of the FHIR standard enables the ingestion of diverse data types, including binary formats such as medical imaging (e.g., DICOM), documents, and clinical resources, thereby enhancing the system's capability to process and interpret complex healthcare data.
Security and legislation. Security is vital in digital healthcare systems, which handle sensitive patient data. Ensuring security in digital healthcare platforms requires strict adherence to the Confidentiality, Integrity, and Availability (CIA) principles while complying with regulatory frameworks to protect sensitive patient data. Confidentiality ensures that only authorized users can access protected data, while integrity guarantees protection against unauthorized manipulation during storage or transmission. Availability requires timely and uninterrupted access to data and services. Security and legal aspects are closely linked to data protection and privacy regulations. In the EU, the General Data Protection Regulation (GDPR) establishes strict guidelines for managing user and health information. *Confidentiality* is ensured through robust authentication and authorization mechanisms, including OAuth2.0, OpenID Connect (OIDC), and Multi-Factor Authentication (MFA). Access control is enforced via Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), ensuring that only authorized users can access or modify data. The SMART-on-FHIR framework further strengthens security by integrating authentication and authorization mechanisms tailored to healthcare interoperability. *Integrity* is maintained by implementing encryption techniques at different levels. *Application-level encryption* secures data stored in databases, file system encryption protects stored information from unauthorized access, and TLS over HTTPS ensures s*ecure data transmission*, preventing threats such as eavesdropping and man-in-the-middle attacks. *Availability* is achieved through redundancy strategies, system-wide monitoring, and disaster recovery plans, ensuring continuous access to critical healthcare services. Intrusion detection systems, anomaly tracking, and auditing mechanisms (e.g., FHIR's Audit resource) further enhance system reliability and compliance. Security and *legal compliance* are closely intertwined. Regulations such as the General Data Protection Regulation (GDPR) in the EU establish strict requirements for handling personal health data. Platforms must implement user consent management, ensuring explicit, revocable, and well-documented patient consent. Furthermore, anonymization and pseudonymization techniques help protect patient privacy when data is used for analytics and external

processing. Regulatory frameworks also impose obligations for cyberattack notification (e.g., within 72 hours under GDPR) and require comprehensive auditing and logging to track access and modifications.

Customization, personalization, and self-adaptation. The digital healthcare platform must be inherently *customizable* to adapt to the specific requirements of different monitoring scenarios. Given the variability of medical conditions and monitoring objectives, the system must provide the flexibility to configure treatment plans, select the relevant vital parameters to be tracked, and define the data processing methods necessary for clinical decision-making. This adaptability ensures that the platform can be effectively deployed across diverse healthcare contexts, addressing the unique needs of different patient groups and medical specializations. Beyond clinical adaptability, the platform must also support *personalization* at the individual level. Each patient may require a distinct approach to monitoring and intervention based on their medical history, treatment phase, and overall health condition. The system should allow dynamic configuration of monitoring protocols, alert thresholds, and decision-support features to ensure a patient-specific approach. Moreover, the platform should *self-adapt* based on the operational environment. This means it should adjust to changes in patient conditions, treatment progress, and external factors such as device availability or connectivity constraints.

Performance. Existing solutions often overlook system performance, a crucial factor in large-scale telemedicine platforms. While addressing specific challenges, they neglect scalability and efficiency. Given the platform's diverse functionalities, a distributed architecture is essential, ensuring separation of responsibilities and preventing system overload. A microservices-based design, supported by an API gateway, enables dynamic load balancing and scalable resource allocation. Additionally, clinical data redundancy must be managed effectively. Integration with multiple sources often leads to duplicated or unstructured data, hindering decision-making. The system should eliminate redundancies, store only clinically relevant data, and use temporary storage for intermediate processing, optimizing computational efficiency and storage.

4. The MIRACLE Architecture

The MIRACLE architecture provides a scalable, secure, and interoperable framework for digital health platforms. It ensures interoperability through the HL7 FHIR standard, enabling seamless data exchange. Multimodal data management supports structured EHRs, sensor data, and medical imaging via harmonization techniques. Security and legal compliance are ensured through CIA principles, GDPR adherence, encryption, and access control. Customization and adaptability allow dynamic configuration of monitoring and treatment plans, while performance optimization is achieved through a microservices-based design, ensuring scalability and efficient data processing. MIRACLE delivers a comprehensive, patient-centered solution for modern healthcare needs. As illustrated in Figure 1, the architecture follows a microservice-based design, ensuring clear separation of responsibilities, optimized performance, and horizontal scalability. Below is the description of the key services and their respective roles.

Identity and Access Management (IAM). This service manages user identities and enforces access control policies. It uses the SMART on FHIR standard for secure access to health data according to defined roles and scopes. By decoupling identity management,

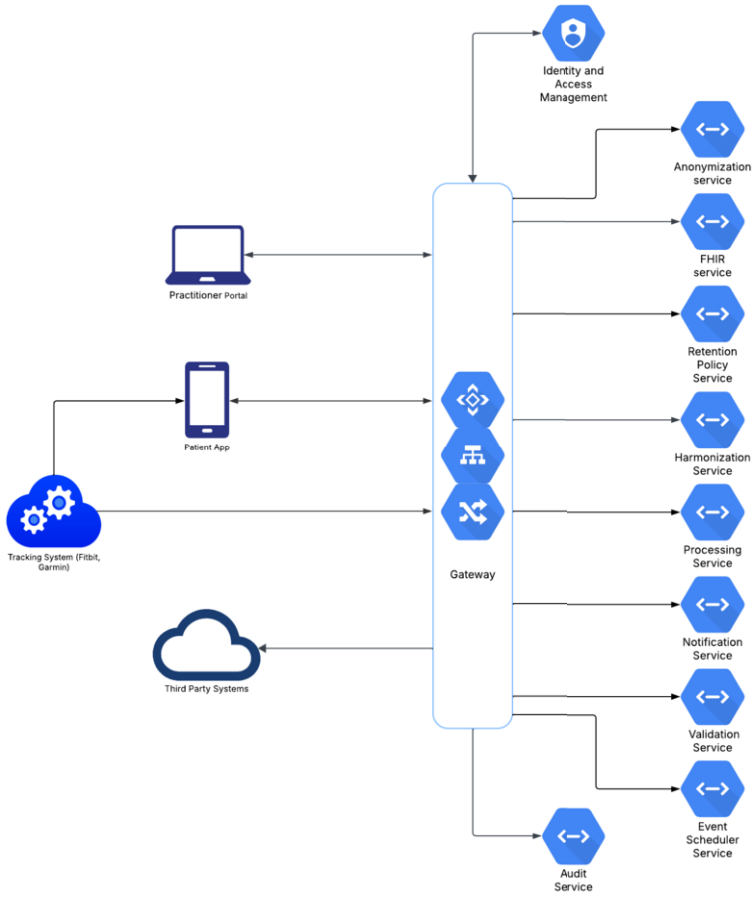

Figure 1. Overview of the proposed architecture, emphasizing its core services.

the system facilitates integration with new security requirements, ensuring compliance with regulations such as GDPR and HIPAA. IAM enforces fine-grained access control and enhances data protection and privacy.

Anonymization Service. This service anonymizes or pseudonymizes patient data, ensuring compliance with privacy regulations. The challenge lies in maintaining data integrity and usefulness for analysis while protecting patient privacy.

FHIR Service. This service serves as the backbone of clinical data management. It is responsible for the permanent storage of health data, ensuring compliance with the FHIR standard to facilitate interoperability across various healthcare systems. The service provides standardized access to patient records, medical histories, and other relevant clinical data. Furthermore, it ensure proper management of patient consent to comply with legal requirements such as GDPR, adding complexity to its responsibilities.

Retention Policy Service. This service manages data retention within the system, ensuring clinical data is stored permanently or temporarily based on medical and regulatory requirements. It guarantees compliance with legal and clinical guidelines while minimizing system load by avoiding unnecessary data retention. The main challenge lies in

determining which data should be kept and ensuring relevant data is preserved, while temporary data is either processed or discarded to avoid performance overhead.

Harmonization Service. This service ingests, structures, and standardizes health data from various sources, ensuring interoperability by converting heterogeneous data into a unified FHIR format. The challenges include managing diverse data types (e.g., text, images, and numerical data) and ensuring proper semantic code mapping for accurate data interpretation.

Processing Service. This service manages temporary data storage and clinical information processing, performing tasks such as data aggregation, analysis, and generating clinical insights for healthcare providers. It is essential to ensure efficient storage without compromising processing capabilities, while defining clear data processing pipelines based on clinical and medical needs.

Notification Service. This service manages multi-channel notifications (e.g., SMS, email, push notifications, in-app messages) to keep users informed about critical events such as appointments or care plan updates. Timeliness and accuracy are essential, as delays in communication can adversely impact patient care.

Validation Service. This service ensures data integrity within the system by validating both incoming FHIR and harmonized data against the required standards. This process is vital for maintaining system reliability, as invalid data could lead to non-compliance with FHIR standards and potentially erroneous clinical insights.

Event Scheduler Service. This service is responsible for managing the timely execution of system and application events, including periodic tasks (e.g., data processing) and one-off events (e.g., appointment notifications). It ensures proper workflow and task coordination, addressing challenges like high workloads and task conflicts. The service must scale effectively to accommodate increasing event volumes.

Audit Service. This service tracks relevant system events, ensuring traceability for compliance with privacy and security regulations. It monitors system performance through key metrics and health checks. A key challenge is minimizing performance overhead while processing large volumes of clinical data.

API Gateway. The API Gateway serves as the central entry point within the architecture. It manages all incoming requests and orchestrates services to ensure interoperability, as well as meet operational and security requirements. Additionally, it facilitates system adaptability to accommodate future changes. As a critical component, the gateway is responsible for routing traffic, performing load balancing, and directing requests to specific services. It also forwards authentication and authorization requests to the Identity and Access Management (IAM) system, ensuring that only authenticated and authorized users can access the system's resources. Given the gateway's pivotal role in facilitating service interactions, particular attention must be paid to avoid creating a Single Point of Failure (SPOF) and bottlenecks. Its scalable design allows for horizontal replication, improving system capacity and resilience.

5. Case Study

This case study focuses on monitoring patients undergoing cardiac surgery, ensuring continuous support throughout their medical journey. The patient's trajectory is divided into several phases, each presenting unique challenges and requirements.

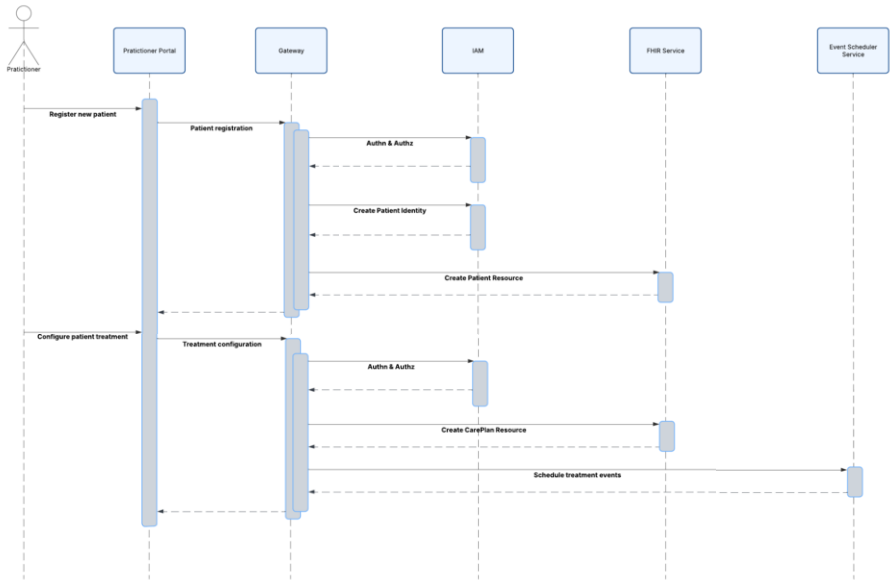

Figure 2. Sequence diagram of new patient registration and treatment configuration.

In the *preoperative phase*, patients diagnosed with cardiovascular conditions requiring surgery are registered in the system. This initiates structured tracking of their medical data and facilitates prehabilitation exercises and lifestyle changes to optimize surgical outcomes. A well-defined system is critical in guiding patients through this phase to minimize post-surgical complications.

The *hospitalization phase* follows, where the patient undergoes surgery and subsequent monitoring of vital signs. After surgery, the patient moves from intensive care to a general ward, where rehabilitation and continuous monitoring are essential for ensuring stability before progressing to the next phase.

The *rehabilitation phase*, often in a specialized clinic, involves remote monitoring, as the primary healthcare provider may no longer have direct contact with the patient. Digital systems that enable remote tracking of recovery indicators can significantly enhance outcomes and facilitate timely medical interventions.

Once discharged, the patient enters the *home-based postoperative phase*, with periodic medical reviews and long-term recovery plans. A remote monitoring system can track health parameters, issue reminders, and provide medical suggestions, thus offering valuable support throughout the recovery process.

Given the complexity of this multi-phase journey, a structured approach to monitoring, data processing, and decision-making is essential at each stage. The proposed architecture ensures seamless integration of monitoring features for both patients and healthcare professionals.

As illustrated in Figure 2, when a patient is registered, their identity is created within the IAM system, and their medical data is securely stored in the FHIR service. At this stage, necessary monitoring parameters, such as medical appointments and assessments, are configured. The Event Scheduler Service then triggers events based on this data.

During active monitoring, the patient interacts with the system through a mobile application, transmitting health data from fitness trackers or manually entered information. This data undergoes an ingestion process that validates its FHIR compliance, evaluates retention policies, and stores it appropriately, either permanently in the FHIR service or temporarily in the Processing Service. Data from external systems is first harmonized by the Harmonization Service to conform to the FHIR standard, as shown in Figure 3.

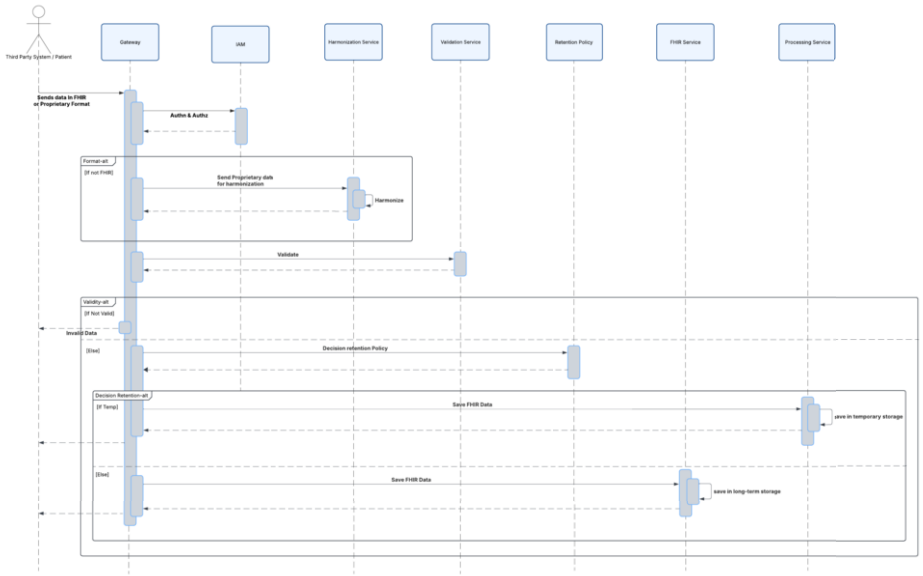

Figure 3. Sequence diagram of data ingestion process.

Once the Event Scheduler triggers relevant processing tasks, the system performs predefined operations for the patient. Results are stored in the FHIR service, and if clinical intervention is required, the Notification Service sends timely updates to both healthcare providers and patients via various communication channels.

We implemented a prototype to evaluate feasibility and effectiveness. The following describes the technologies adopted.

The open source Traefik[1] has been identified as the API gateway. It is a modern reverse proxy and load balancer that automatically discovers services when deployed in environments such as Docker and Kubernetes. Supporting TLS, routing, and middlewares, Traefik offers a customizable and extensible gateway based on architectural needs. It integrates seamlessly with Keycloak[2] as the Identity and Access Management (IAM) solution. Keycloak is an open-source IAM platform that supports authentication, authorization, and single sign-on (SSO) for applications using protocols such as OAuth2, OpenID Connect, and SAML.

[1] https://traefik.io/traefik/
[2] https://www.keycloak.org/

For the FHIR service, Medplum[3], an open-source healthcare API platform, was used to simplify development. This ensured compliance with the FHIR healthcare standard and facilitated seamless data exchange within healthcare systems.

For monitoring, Prometheus[4] and Elasticsearch[5] were selected. Prometheus is a scalable tool for time-series data collection, storing metrics from applications and infrastructure for analysis and alerting. It is ideal for tracking component health. Elasticsearch, with its distributed search engine, is suited for storing and quickly searching large datasets, particularly for complex logging and auditing information.

The processing service uses custom algorithms, either suggested by practitioners or developed ad hoc. For the Notification service, we are exploring multi-channel notification, with Novu[6] as an open-source platform to orchestrate updates via email, SMS, push notifications, and chat.

Patient data, if not already in FHIR, must be converted with syntactic and semantic validation. The data is then stored either temporarily in the Processing Service or permanently in the FHIR service, according to retention policies. Due to the variability in third-party tracker (TPT) data sources, creating custom Harmonization Services for each would be costly. Instead, a single service leveraging Large Language Models (LLMs) offers a more efficient solution, though it may introduce inference latency and potential harmonization errors. Syntactic validation is handled by the official FHIR validator, while semantic validation requires mapping to SNOMED and LOINC codes, potentially aided by a Retrieval-Augmented Generation (RAG) model. Storage decisions will follow physician guidelines or an intelligent custom policy.

To ensure GDPR compliance and facilitate data sharing, we have implemented open-source solutions for data anonymization, including Microsoft FHIR Anonymizer[7], ARX Data Anonymization Tool[8], and OpenDP[9]. For data security, all traffic is securely routed via HTTPS, with certificate management handled by the API gateway using services like Vault[10] or Let's Encrypt[11]. Data at rest is encrypted using Trusted Secure Databases (TSD) or volume-level encryption, while PostgreSQL Transparent Data Encryption (TDE) is applied for relational databases.

6. Conclusions

The MIRACLE architecture delivers a modular, secure, and interoperable framework designed to address the multifaceted needs of remote patient monitoring. By leveraging HL7 FHIR standards, microservices, and privacy-preserving mechanisms, the architecture supports the integration of heterogeneous data sources, adaptability to diverse clinical settings, and compliance with GDPR and HIPAA regulations. Its services are clearly defined and organized to ensure scalability, performance optimization, and fine-

[3] https://www.medplum.com/
[4] https://prometheus.io/
[5] https://www.elastic.co/elasticsearch
[6] https://novu.co/
[7] https://github.com/microsoft/Tools-for-Health-Data-Anonymization
[8] https://arx.deidentifier.org/
[9] https://opendp.org/
[10] https://www.vaultproject.io/
[11] https://letsencrypt.org/

grained control over identity, consent, and data retention. We recognize that adopting a microservice-based design introduces architectural complexity and that implementing stringent security measures may impact system performance. Nonetheless, the modularity enabled by microservices significantly enhances system adaptability and scalability. At the same time, ensuring security and privacy is essential to maintain user trust and enable sustainable remote patient monitoring.

While the prototype demonstrates feasibility in cardiac surgery follow-up, future work will assess performance under varying workloads, integration with legacy systems, and microservice orchestration. We also aim to deepen the security analysis to explore robustness-efficiency trade-offs.

Acknowledgment

This work was partially funded by the National Plan for NRRP Complementary Investments (PNC, established with the decree-law 6 May 2021, n. 59, converted by law n. 101 of 2021) in the call for the funding of research initiatives for technologies and innovative trajectories in the health and care sectors (Directorial Decree n. 931 of 06-06-2022) - project n. PNC0000003 - AdvaNced Technologies for Human-centrEd Medicine (project acronym: ANTHEM).

References

[1] Hornback A, Shi W, Giuste FO, Zhu Y, Carpenter AM, Hilton C, et al. Development of a generalizable multi-site and multi-modality clinical data cloud infrastructure for pediatric patient care. In: Proceedings of the 13th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics. BCB '22. New York, NY, USA: Association for Computing Machinery; 2022. Available from: https://doi.org/10.1145/3535508.3545565.

[2] Marfoglia A, Nardini F, Arcobelli VA, Moscato S, Mellone S, Carbonaro A. Towards real-world clinical data standardization: A modular FHIR-driven transformation pipeline to enhance semantic interoperability in healthcare. Computers in Biology and Medicine. 2025;187:109745. Available from: https://www.sciencedirect.com/science/article/pii/S0010482525000952.

[3] Li Y, Wang H, Yerebakan H, Shinagawa Y, Luo Y. Enhancing Health Data Interoperability with Large Language Models: A FHIR Study; 2023. Available from: https://arxiv.org/abs/2310.12989.

[4] Rindal DB, Pasumarthi DP, Thirumalai V, Truitt AR, Asche SE, Worley DC, et al. Clinical Decision Support to Reduce Opioid Prescriptions for Dental Extractions using SMART on FHIR: Implementation Report. JMIR Med Inform. 2023 Nov;11:e45636.

[5] dos Santos Leandro G, Moro CMC, Cruz-Correia RJ, Portela Santos EA. FHIR Implementation Guide for Stroke: A dual focus on the patient's clinical pathway and value-based healthcare. International Journal of Medical Informatics. 2024;190:105525. Available from: https://www.sciencedirect.com/science/article/pii/S1386505624001886.

[6] Alamri B, Javed IT, Margaria T. A GDPR-Compliant Framework for IoT-Based Personal Health Records Using Blockchain. In: 2021 11th IFIP International Conference on New Technologies, Mobility and Security (NTMS); 2021. p. 1-5.

[7] Chatterjee A, Pahari N, Prinz A. HL7 FHIR with SNOMED-CT to Achieve Semantic and Structural Interoperability in Personal Health Data: A Proof-of-Concept Study. Sensors. 2022;22(10). Available from: https://www.mdpi.com/1424-8220/22/10/3756.

[8] Vasileiou N, Giannakopoulou O, Manta O, Bromis K, Vagenas TP, Kouris I, et al. FHIR-Driven Advancements in Healthcare Interoperability: Insights from the Retention Project. In: 2024 IEEE International Conference on Engineering, Technology, and Innovation (ICE/ITMC); 2024. p. 1-6.

[9] Raso E, Loreti P, Ravaziol M, Bracciale L. Anonymization and Pseudonymization of FHIR Resources for Secondary Use of Healthcare Data. IEEE Access. 2024;12:44929-39.