

Article

Not peer-reviewed version

PIRE: Interoperable Platform for Electronic Records

[Leonardo Juan Ramirez Lopez](#)*, Norman Eduardo Jaimes Salazar, [Juan Esteban Barbosa Posada](#)

Posted Date: 28 January 2026

doi: 10.20944/preprints202601.2177.v1

Keywords: HL7 FHIR; semantic interoperability; biosignals; data government



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

PIRE: Interoperable Platform for Electronic Records

Leonardo Juan Ramirez Lopez *, Norman Eduardo Jaimes Salazar
and Juan Esteban Barbosa Posada

TIGUM Research Group, Universidad Militar Nueva Granada, Bogota 110111, Colombia

* Correspondence: leonardo.ramirez@unimilitar.edu.co; Tel.: +57-1-6500000

Abstract

The interoperability of electronic health records in Colombia faces a critical gap between the regulatory mandates established by Law 2015 of 2020, Resolution 866 of 2021, and Resolution 1888 of 2025, and the actual technical capacity of healthcare institutions to implement them. This paper presents PIRE (Electronic Records Interoperability Platform), an open-source architecture that demonstrates the viability of end-to-end FHIR systems in the Colombian context. The main objective was to develop a platform capable of integrating health data from biomedical devices into an FHIR server, preserving clinical semantics through LOINC terminologies. The methodology followed an iterative development approach, implementing a HAPI FHIR server on AWS, a normalization application in Flask, and clinical visualization modules aligned with the FHIR Core CO Implementation Guide. The Bioharness-3 device was used to capture metrics on heart rate, respiratory rate, activity, and posture. The results demonstrate that the architecture enables the semantically preserved exchange of biosignals in real time, validating compliance with Colombian IHCE specifications. It is concluded that PIRE constitutes a reproducible reference model for healthcare institutions that wish to implement interoperability without relying on costly enterprise solutions.

Keywords: HL7 FHIR; semantic interoperability; biosignals; data government

1. Introduction

1.1. Electronic Health Record

The digital transformation of the healthcare sector is based on the premise that clinical information, when available in a structured, integrated, and accessible form, can significantly improve the quality and safety of healthcare. It is in this context, electronic health records (EHR) have emerged as a response to the fundamental need of modern healthcare systems. An electronic health record is a digital information system that captures, stores, organizes, and facilitates access to a patient's medical records in an integrated manner over time and across multiple healthcare encounters. The evolution of the EHR has unfolded through milestones over time. Shen et al. place the conceptual origins of EHR in the 1960s, highlighting milestones such as the Problem-Oriented Medical Record (POMR) introduced by Lawrence Weed in 1968 and the development of the Regenstrief Medical Record System (RMRS) developed by Clement McDonald in 1972-1973, which laid the foundation for electronic data capture and retrieval [1]. Conversely, Finnegan and Mountford point out that, over the past 25 years, EHRs have evolved from rudimentary digital repositories to tools that optimize patient safety, driven by technological innovations and fundamental reforms [2]. More specifically, Shen et al. situate this recent progression by dividing it into decades, where the 2000s focused on individual hospital systems and image digitization, the 2010s saw the rise of Big Data and real-world evidence, and the 2020s are marked by the advanced use of artificial intelligence (AI), genomics, and real-time public health surveillance. [1]. This evolution has resulted in widespread adoption in modern healthcare systems. According to data from the Office of the National Coordinator for Health Information Technology (ONC), 90% of hospitals in the United

States operate with EHR systems, 99% allow electronic access to records by patients, and 70% participate in the four domains of interoperability (send, receive, search, integrate) [3]. Similarly, Shen et al. document contemporary global-scale implementations such as UK Biobank with over 500,000 participants and All of Us with over 760,000 participants in the United States, which integrate genomic data for precision medicine [1]. However, the fragmentation of EHR systems built by different providers with heterogeneous data structures and operating under diverse national regulations has created a new critical challenge that dominates the current research agenda. Specifically, El-Yafouri and Klieb document that, although technical and syntactic levels of interoperability have reached relative maturity, semantic interoperability—the ability to ensure that clinical information is understood identically when exchanged between systems—remains insufficiently resolved [4]. As a result, this fragmentation results in institutionally segregated patient records, impeding continuity of care, integrated research, and effective public health surveillance.

1.2. *The Interoperability of Electronic Health Records*

The interoperability of health records is the backbone that allows patient information to flow securely between systems, institutions, and levels of care. Beyond being merely a technical issue, interoperability encompasses different layers, from the basic ability to transmit data to ensuring that this data retains its clinical meaning and fits into real-world processes. Various experiences illustrate how countries have addressed this need using legacy standards. In Austria, Duftshmid et al. analyze the shared EHR system ELGA, based on HL7 Clinical Document Architecture (CDA), to assess whether its content can be reused in conducting clinical trials. Using semi-automatic mapping with the ART.DECOR tool, the authors show that ELGA covers 88% of relevant data elements and that 77% of these are available in structured format, demonstrating a high degree of syntactic interoperability and considerable potential for secondary data reuse [5]. Complementarily, Palojoki et al. conduct a systematic review of 14 implementations of semantic interoperability in EHRs, identifying the prevalent use of ontologies, clinical terminologies such as SNOMED CT and LOINC, and standards such as openEHR, CDA, and FHIR. The analysis confirms that preserving clinical meaning during data exchange between heterogeneous systems is the primary goal of these efforts [6].

Nevertheless, the complexity involved in maintaining architectures based exclusively on legacy standards such as HL7 CDA has revealed limitations in the face of growing demand for dynamic exchange and advanced data reuse. To address this transition, Bossenko et al. address this transition through TermX, an open-source platform that implements FHIR Mapping Language (FML) to transform CDA documents into FHIR resources using reusable visual components. When evaluated in the context of the European Health Data Ecosystem (EHDS), this solution demonstrates that clinical experts can define semantically consistent transformations between heterogeneous national systems, facilitating the progressive migration to FHIR architectures [7].

These technical developments are situated within a global regulatory context that has elevated interoperability to an unavoidable regulatory requirement. The World Health Organization (WHO) documents that, in the European Region, the advancement of shared EHR systems depends on national architectures aligned with international standards, particularly within the framework of the European Health Data Space, which promotes standardized FHIR profiles [8]. In the United States, the Office of National Coordination for Health Information Technology (ONC) reports that 70% of hospitals are actively participating in the four domains of interoperability, driven by regulations that require FHIR-based APIs and minimum data sets such as USCDI [3]. Notably, at the same time, international surveys coordinated by HL7 International reveal that, by 2025, 71% of countries report active adoption of FHIR for specific clinical use cases, with 73% of national regulations explicitly requiring or recommending its implementation [9]. Taken together, the convergence of technical, regulatory, and global adoption evidence positions the FHIR standard as the dominant response to persistent semantic interoperability challenges in modern EHR systems, particularly in contexts of multidisciplinary care and secondary data reuse.

1.3. Fast Healthcare Interoperability Resources

FHIR is a standard designed to provide a single set of resources that serve as building blocks for interoperability solutions in healthcare contexts. According to Andrew Hornback et al., the FHIR standard has established itself as the predominant technology for harmonizing biomedical data, facilitating the integration of electronic records, predictive analytics, and decision support tools [10]. Importantly, this digital transformation is not limited to advanced systems, as Abar Ahmad et al. demonstrate that it is feasible to implement FHIR even in developing countries using the SFPBRF framework, which allows paper-based records to be mapped to standardized resources to improve patient care [11]. In the field of clinical research, Alex C. Cheng et al. emphasize that integrating tools such as REDCap with electronic records via FHIR drastically reduces human error, although organizational barriers and the need for IT staff training remain [12]. To facilitate these technical transitions, Jesse Kruse et al. introduce FML2Mirth, a solution that automates the generation of integration channels using the FHIR Mapping Language (FML) to process complex laboratory data [13]. Similarly, Mie Vestergaard Andersen et al. confirm the feasibility of migrating complex legacy models, such as Danish microbiology systems, to the FHIR standard, although they caution that this requires intensive use of extensions to capture local particularities [14]. In parallel, Maryna Khvastova et al. propose the use of plugins to provide FHIR interoperability to XNAT open-source research platforms, enabling the seamless exchange of study subject data [15].

The convergence of FHIR with advanced artificial intelligence is explored by Merlin Engelke et al., who present FHIR-Former, a system that uses language models (LLMs) to process multimodal data and predict risks of mortality or hospital readmission [16]. Building on this, Zhen Hou et al. introduce LINK-FHIR, a tool that leverages the reasoning capabilities of models such as GPT-4 to standardize unstructured clinical notes and laboratory reports into FHIR resources with high accuracy [17]. For the massive-scale analysis of this data, John Grimes et al. detail the SQL on FHIR standard, which projects nested information into tabular views, allowing data analysts to use traditional tools to replicate complex clinical studies in a portable manner [18].

Regarding security and governance, Subnojeet Mukherjee et al. propose an attribute-based access control (ABAC) model that allows granular policies to be defined regarding who can view or edit specific resources [19]. At the distributed security level, Antonio López Martínez et al. present an arc that uses Sovereign Identity (SSI) and Blockchain to manage access permissions to FHIR resources, including emergency protocols for unconscious patients [20]. Equally important is usability, and Carsten Vogel et al. evaluate FHIR visual questionnaire editors using eye tracking to ensure that both physicians and developers can manage data without cognitive errors [21].

The standard's specialization extends to critical domains. In genomic medicine, Nina Haffer et al. demonstrate the feasibility of mapping national datasets to international genomic reporting guidelines [22]. In hospital operations management, Cynthia Sabrina Schmidt et al. implemented a real-time dashboard to optimize the use of blood products, identifying units reserved for patients who had already been discharged [23]. In the biopharmaceutical field, Craig Anderson et al. propose the use of FHIR for the exchange of pharmaceutical product quality data (CMC), accelerating regulatory processes from years to days [24]. Additionally, Karamarie Fecho et al. present FHIR PIT, an integration pipeline that links clinical records with geospatial data on environmental pollution to study impacts on diseases such as asthma [25]. For the biobank sector, Antonella Cruoglio et al. detail an ETL process that converts minimal metadata into FHIR transaction packages, enabling federated searching of biological samples across European networks [26]. Finally, Junghoon Lee et al. improve anomaly detection in ECG data streams using a voting scheme that reduces false negatives in remote monitoring [27].

With respect to real-time monitoring and mobile devices, Adil Akhmetev et al. present a platform that combines FHIR with Federated Learning (FL) to train predictive models on physical activity while preserving the privacy of wearable sensor data [28]. Similarly, Fabio A. Seixas-Lopes et al. apply FHIR for the personalized collection of data on musculoskeletal disorders, integrating posture and activity metrics through an Internet of Medical Things (IoMT) architecture [29]. This

patient-centered approach is shared by Pietro Randine et al., who developed Dia-Continua, a system that uses patient-generated data for type 1 diabetes consultations, linking glucose levels and insulin doses under the FHIR standard [30]. In related vein, Philipp Urbauer et al. validate the use of FHIR to integrate activity trackers into elderly care, enabling semantic monitoring of sleep patterns and heart rate [31]. Similarly, Somayeh Abedian et al. highlight how the integration of commercial devices such as Garmin into the European Health Data Space (EHDS) facilitates clinical decision-making based on continuous health data [32]. To provide technical support for these networks, Felix-Constantin Adochiei et al. propose an open-source framework based on ESP32 microcontrollers and the MQTT protocol to transmit biomedical signals such as ECG and PPG directly to servers [33].

In summary, FHIR's evolution from its adoption in traditional EHR systems to its integration with emerging technologies such as artificial intelligence, federated learning, and multiparametric wearable devices demonstrates the standard's maturity in orchestrating complex and heterogeneous digital health ecosystems. This convergence of capabilities, semantic interoperability, advanced predictive analytics, granular security, and continuous patient data collection forms the technological foundation upon which it is possible to implement interoperable solutions that integrate continuous biomedical monitoring devices into real clinical workflows. However, the viability of these solutions in specific contexts such as Colombia depends fundamentally on how they are articulated with the regulatory frameworks, technological architectures, and interoperability standards currently deployed in national health institutions, particularly with regard to the integration of new sources of biomedical data with existing electronic health records.

1.4. Electronic Health Records in Colombia

In Colombia, the interoperability of Electronic Health Records has gone from being an aspiration to a legal obligation, based on a regulatory framework that articulates Law 2015 of 2020, Resolution 866 of 2021, and Resolution 1888 of 2025, which converges on the use of HL7 FHIR as the technical reference standard for Interoperable Electronic Health Records (IHCE). Specifically, Law 2015 establishes the IHCE as a mandatory national system, mandating the Ministry of Health to define standards and exchange mechanisms to ensure continuity of care, patient safety, and timely access to clinical data [34]. Build on this, Resolution 866 specifies the set of relevant clinical data that must be structured and transmitted uniformly by health service providers [35]. Most recently, Resolution 1888 of 2025 expands on this framework by adopting the Digital Care Summary (RDA) as the centerpiece of the information flow, regulating how local health centers must generate, send, and make available this summary through the national interoperability platform, thereby closing the circle between clinical records, data exchange, and secondary use of information [36]. To operationalize this legal framework supports, technical instruments such as the Vulcano national reference server and the FHIR Core Co Implementation Guide, which provide profiles, extensions, terminologies, and specific examples of FHIR resources (e.g., Patient-CO, PractitionerCO, EncounterCO, and associated clinical resources) adapted to the Colombian regulatory and operational context, allowing developers and institutions to align their implementations with the IHCE/RDA model and with the identification, coding, and security requirements defined by the Ministry of Health [37].

In line this regulatory framework, some actors in the Colombian healthcare system have made progress in implementing specific interoperability solutions based on HL7 FHIR, which serve as a practical reference for IHCE in production. In the case of SURA, the insurer undertook a seven-year technology evaluation process that included early implementation of HL7 FHIR DSTU2 in 2016, culminating in 2023 with the design of an interoperability architecture on Google Cloud Healthcare API, adopting HL7 FHIR R4 and developing its own guide with more than fifty FHIR profiles aligned with the Colombian CORE specification. This platform, deployed in less than twelve months and designed to be integrated with the national RDA in the future, now enables the structured exchange of clinical information between multiple points of care within the network [38]. In the corporate sphere, Ecopetrol has promoted a certified integration model that requires contracted IPSs to interoperate using HL7 FHIR for the exchange of administrative and clinical data. Through solutions

such as Medifolios, automatic two-way synchronization of demographic data, affiliations, authorizations, and coverage status is established, along with the sending and receiving of diagnoses, procedures, medications, and medical progress reports in FHIR format, also incorporating real-time rights validation and auditable traceability of each transaction [39,40].

The implementation of pilot tests and integration of biomedical devices using the FHIR standard has advanced significantly. Nevertheless, the set of regulatory advances and implementation cases demonstrates that the Colombian ecosystem is sufficiently mature to experiment with new data sources and clinical architectures. Critically, in Colombia, academic evidence documenting pilot tests of FHIR architectures remains very limited, with published studies focusing on public policy analysis and adoption diagnostics, but rarely describing in detail actual implementations with FHIR servers, integration with legacy systems, or medical devices, especially in university settings. This gap between regulatory design and the production of technical evidence on interoperable architectures hinders the consolidation of implementations aligned with Colombian regulations and scientific advances in clinical interoperability.

1.5. Highlights

This study develops and evaluates the Interoperable Electronic Records Platform (PIRE) that maps PGHD data from the Bioharness-3 device (Zephyr Technology™) device to HL7 FHIR R4 resources aligned with the Colombian RDA model (Resolution 1888/2025), deployed on a HAPI FHIR server on AWS EC2 and validated in pilot tests at the Universidad Militar Nueva Granada. The integration of continuous biomedical signals (heart rate, respiratory rate, activity, and posture) into standardized clinical workflows is demonstrated, generating dynamic PatientCO and PractitionerCO profiles that comply with FHIR Core CO specifications, managing a system of roles and permissions for access to platform information, and supporting secondary clinical-predictive analyses. Collectively, these contributions address the gap between national regulations and providing technical evidence on PGHD data interoperability in Colombian university contexts.

2. Materials and Methods

2.1. Materials

The Interoperable Electronic Records Platform (PIRE) is an open, HL7 FHIR R4-based interoperability architecture that includes web services and tools for the comprehensive management of clinical and biomedical data. To achieve this, PIRE was deployed on Amazon Web Services (AWS) cloud infrastructure, using an EC2 compute instance running Ubuntu Linux as the runtime environment for the HAPI FHIR server and Flask web application. This configuration allows the platform to be remotely accessible to authorized users at the Universidad Militar Nueva Granada, while simultaneously facilitating centralized administration of clinical resources and future scalability of the system.

2.1.1. Technology Stack

The platform's technology stack combines Python web development components with Java-based FHIR server infrastructure and related persistence via PostgreSQL. Specifically, the web application layer was implemented using Python's Flask framework, which manages HTTP routes, data capture forms, authentication and authorization logic, and communication with the FHIR server via REST requests. To enhance security, Flask is complemented by Flask-Session for server-side session management, allowing you to securely store authentication status and temporary data from observation log streams. In parallel, the HAPI FHIR server runs on the Java Spring Boot framework, using Open JDK 17 as the runtime environment and Maven as the build and dependency management tool. For persistence purposes, FHIR resource persistence is achieved using PostgreSQL 16, which stores Patient, Observation, and Practitioner resources in JSON formats within specialized

tables, thereby enabling direct SQL queries on clinical data when advanced analysis or validation of stored information is required. The communication pathway between the Flask application and the HAPI FHIR server is carried out using the standard FHIR RESTful API, using HTTP GET, POST, and DELETE operations on the endpoints corresponding to each type of resource. Table 1 below presents details about the PIRE technology stack.

Table 1. PIRE technology stack.

Layer	Component	Technology	Version
Web application	Framework	Flask (Python)	2.x
Web application	Sessions	Flask-Session	-
FHIR server	Implementation	HAPI FHIR JPA Server	8.2.0
FHIR server	Runtime	OpenJDK	17.0.15
FHIR server	Build	Maven	3.x
Database	Engine	PostgreSQL	16.9
Operating system	Distribution	Ubuntu Linux	24.04.2 LTS
Infrastructure	Cloud	AWS EC2	-

2.1.2. HAPI FHIR R4 Server

The HAPI FHIR server was installed from the official `ha-pi-fhir-jpaserver-starter` repository, which provides a production-ready implementation of the FHIR R4 standard with support for JPA persistence and a test web interface. The installation process began with cloning the repository using Git, followed by installing the Java and Maven environment needed to compile the project, using standard Ubuntu package management commands. Subsequently, the connection to the PostgreSQL database was configured by specifying the connection URL and the database administrator's username and password. Once configured, the server is executed using the Maven Spring Boot plugin, which starts the application and exposes the FHIR endpoints in the `/fhir` path, thereby enabling the Flask application to consume the RESTful services to create, query, and delete clinical resources. Notably, the HAPI FHIR server stores resources in the PostgreSQL database using a predefined table schema, which contains the metadata for each resource and stores the complete JSON content of each version of the resource, thus permitting both API access and direct SQL queries for validation and analysis.

2.1.3. AWS EC2 Configuration

The computing infrastructure was deployed on an Amazon Web Services EC2 instance, using a virtual machine image with Ubuntu 24.04.2 LTS. To enable access, the instance was configured with a public IP address that allows remote access via SSH, providing a secure mechanism for server administration without exposing user credentials. In practice, access to the server is achieved using the standard SSH command, specifying the path to the private key and the public IP address of the instance, after which the administrator gains access to manage the services. With respect to network configuration, the instance's network configuration includes security group rules that allow incoming traffic on the ports necessary for system operation. Internally, the instance maintains a private IP address within the AWS VPC for internal communication between components, while simultaneously the public IP address allows access from authorized external networks. This deployment scheme facilitates future horizontal scalability, as it allows the instance to be replicated or mitigated to AWS managed services without modifying the application architecture.

2.1.4. Integration with the Universidad Militar Nueva Granada

The platform was developed within the context of the TIGUM research group at the Universidad Militar Nueva Granada, which determined design decisions aimed at facilitating academic use and clinical validation in controlled environments. Specifically, integration with the university's

infrastructure was achieved through remote access to the EC2 server, allowing authorized researchers and students to interact with the platform. To reflect organizational structure, the platform's user and role system was designed to mirror the organizational hierarchy of a clinical-academic environment, with distinct roles for administrators, healthcare professionals involved in research, support staff, and auditors. For data isolation purposes, the PostgreSQL database was configured with a schema called *pire_db* and a specific administrator user for the platform, which allows the platform data to be isolated from other possible uses of the server and simultaneously facilitates the management of backups and eventual migration to permanent institutional infrastructure. To promote knowledge transfer, the technical documentation for the system, including installation guides, architecture, and troubleshooting, was kept in Markdown format within the studio repository, thereby facilitating the transfer of knowledge to future developers or researchers who continue the work. Overall, this configuration allows the platform to function as a working prototype for validating the FHIR interoperability model in the Colombian context, while laying the groundwork for eventual implementation in real clinical production environments.

2.2. Methods

The development of health informatics software adhering to the HL7 FHIR standard requires a multidisciplinary approach that integrates regulatory alignment, data modeling, semantic validation, and user-centered design. Accordingly, PIRE's development followed a structured methodology organized into five sequential phases: (1) regulatory and normative analysis, (2) semantic and functional modeling, (3) architectural design and implementation, (4) security, governance, and user-centered design, and (5) validation and deployment.

2.2.1. Phase 1: Regulatory and Normative Foundation

The development of PIRE began with a comprehensive documentary analysis to establish the normative and regulatory framework governing electronic health records interoperability in Colombia. This phase involved examination of Law 2015 of 2020, Resolution 866 of 2021, and Resolution 1888 of 2025, which collectively establish the Interoperable Electronic Health Record (IHCE) as a mandatory national system. Simultaneously, a systematic review of international FHIR implementations and use cases was conducted to identify successful technical patterns applicable to the Colombian context. This included analysis of implementations in the European Health Data Space (EHDS), the United States Interoperability Standards Advisory, and emerging IoMT-FHIR architectures documented in the peer-reviewed literature. The output of this phase was a comprehensive requirements specification document that articulated both functional and non-functional requirements for PIRE, derived from regulatory mandates and international best practices.

2.2.2. Phase 2: Semantic Data Modeling

To ensure that clinical data retained its precise meaning across system boundaries, PIRE's semantic layer was constructed using standardized clinical terminologies and resource profiles. Specifically, the FHIR R4 Patient, Observation, and Practitioner resources were adapted to comply with the Core Colombia Implementation Guide, incorporating Colombian-specific identifiers, terminology mappings, and organizational hierarchies. For vital signs data, LOINC (Logical Observation Identifiers Names and Codes) terminology was used to standardize the identification of measured phenomena, ensuring that metrics such as heart rate, respiratory rate, and body temperature could be universally understood. Additionally, where applicable, SNOMED CT concepts were mapped to Colombian disease classifications (CIE-10) to preserve clinical semantics during data exchange. This semantic modeling was documented using FHIR StructureDefinitions and implemented as validation schemas within the HAPI FHIR server, ensuring that all resources submitted to the platform conformed to the defined semantic contracts before persistence.

2.2.3. Phase 3: Architectural Design and Implementation

PIRE's architecture was designed following REST (Representational State Transfer) principles, with clinical data organized as persistent, addressable resources accessible via standard HTTPS protocols and HTTP operations (GET, POST, DELETE). This resource-oriented approach contrasts with message-based architectures, as it treats each clinical entity (Patient, Observation, Practitioner) as a first-class resource with a stable URI and standardized access patterns. The technical implementation employed a microservices topology, with distinct components responsible for specific functions: the HAPI FHIR server manages resource persistence and retrieval; the Flask web application provides the user interface and orchestrates business logic; the Bioharness-3 normalization module transforms raw wearable device data into FHIR Observations; and the audit logging subsystem maintains non-repudiation records of all data transactions. Communication between components uses lightweight, language-agnostic formats: JSON for API payloads and structured data interchange, and standard FHIR transaction bundles for atomic multi-resource operations. This modular design facilitates independent scaling of components and reduces coupling between layers.

2.2.4. Phase 4: Security, Governance, and User-Centered Design

Security and governance were implemented through multiple complementary mechanisms. At the authentication layer, role-based access control (RBAC) was employed, with distinct roles (administrator, clinician-researcher, support staff, auditor) corresponding to organizational positions and clinical responsibilities. At the data layer, all access to FHIR resources was logged with comprehensive audit trails recording the timestamp, user identity, operation type, and resource identifiers, enabling post-hoc accountability and compliance with Colombian health data protection regulations (Habeas Data). Additionally, user interface design followed human-centered design principles, prioritizing clarity and cognitive accessibility for clinical users unfamiliar with FHIR's technical complexity. Interface elements were iteratively refined based on feedback from university researchers and students to ensure that visualization of vital signs data and clinical metadata did not introduce cognitive burden or opportunities for error.

2.2.5. Validation Approach

To validate the technical feasibility and functional correctness of PIRE, a pilot study was conducted using the Bioharness-3 wearable device integrated with the platform. Continuous vital signs data (heart rate, respiratory rate, activity, posture) were collected during controlled laboratory sessions, normalized into FHIR Observation resources, and transmitted to the cloud FHIR server via the edge computing normalization module. Success criteria included: (1) zero data loss during transmission, (2) conformance of transmitted resources to FHIR Core Colombia profile specifications, (3) end-to-end latency <600 milliseconds, and (4) successful retrieval and visualization of data through the web application. Additionally, the semantic preservation of measurement metadata (device type, derivation method, temporal precision) was verified by examining the FHIR Observation structure and confirming that all clinical context was retained without loss of meaning.

3. Development

3.1. Operational Flow

This section describes the technical architecture, implementation details, and functional capabilities that enable PIRE to orchestrate four critical functionalities: comprehensive patient management through CRUD operations with FHIR Colombia extensions (PatientCO), intake and processing of biomedical observations from CSV files with automatic signal classification (heart rate, respiratory rate, activity, and posture), clinical consultation and analysis of measurements with automatic stratification of values according to interpretive ranges, and role-based authentication and

authorization systems with comprehensive transaction auditing. As illustrated in Figure 1, the complete operational flow of the platform integrates these components into a cohesive system architecture.

Additionally, the software architecture is described in detail, encompassing component topology, the implementation of role and permission management using security decorators, the generation of FHIR Observation resources with standardized LOINC codes and PatientCO and PractitionerCO profiles aligned with the FHIR Core CO Implementation Guide, and credential management using XOR + Base64 encryption with server-side sessions. Collectively, these elements form an integrated technical foundation that operationalizes the regulatory requirements established by Colombian health informatics standards.

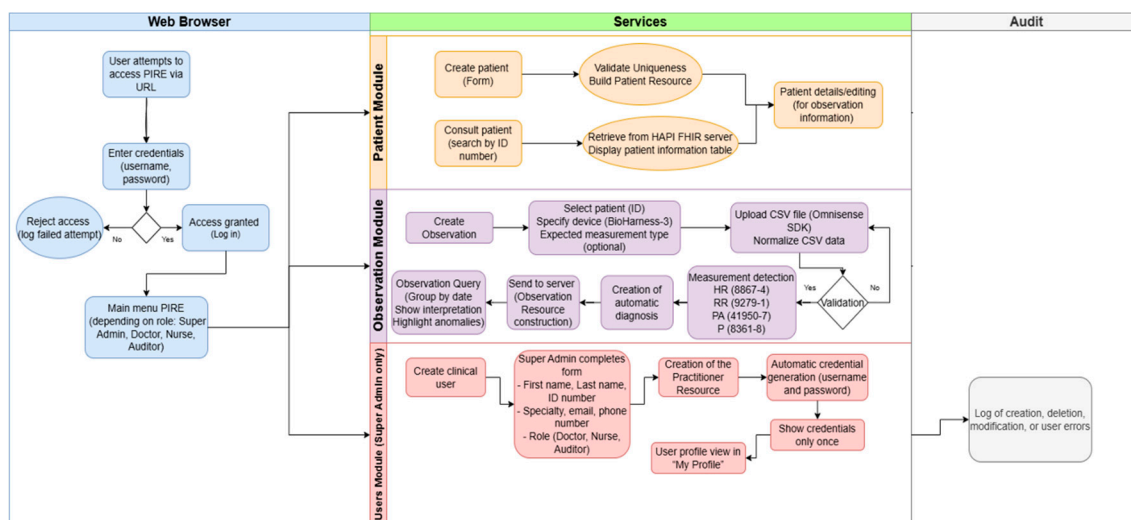


Figure 1. Complete operational flow of PIRE.

3.2. Technical Architecture of the Platform

PIRE adopts a three-layer architecture that separates responsibilities between the user interface, business logic, and data persistence, thereby facilitating system maintainability, scalability, and security. This architectural model allows the Flask application to act as a central orchestrator, integrating specialized modules for authentication, patient management, processing of biomedical observations generated by the Bioharness-3 device, and generation of standardized FHIR R4 resources. Crucially, communication between layers is carried out through RESTful calls to the HAPI FHIR server, ensuring that all clinical data is structured in accordance with international specifications and Colombian regulations from the point of storage.

3.2.1. General System Topology

The PIRE architecture is organized into three functional layers that operate in an integrated manner. At the presentation level, the presentation layer acts as the single interface where users with different roles interact with HTML forms, view clinical information, and upload CSV files containing measurements from the Bioharness-3 device. HTTP requests from the browser are received by the application layer, where routes protected by security decorators validate permissions before processing business logic. Once the user is authenticated and authorized, the authentication module verifies credentials stored in CSV files, while specialized service modules perform operations on patient, observation, and professional data. At the data layer, the data layer comprises two complementary subsystems: a local file system in CSV format that stores user management and audit information, and a HAPI FHIR R4 server that acts as a reference repository for all clinical resources (Patient, Observation, Practitioner) backed by a PostgreSQL database and ensuring compliance with the FHIR R4 standard and alignment with the FHIR Core CO implementation guide. As shown in Figure 2, the data flow architecture presents PIRE's three layers and their interactions.

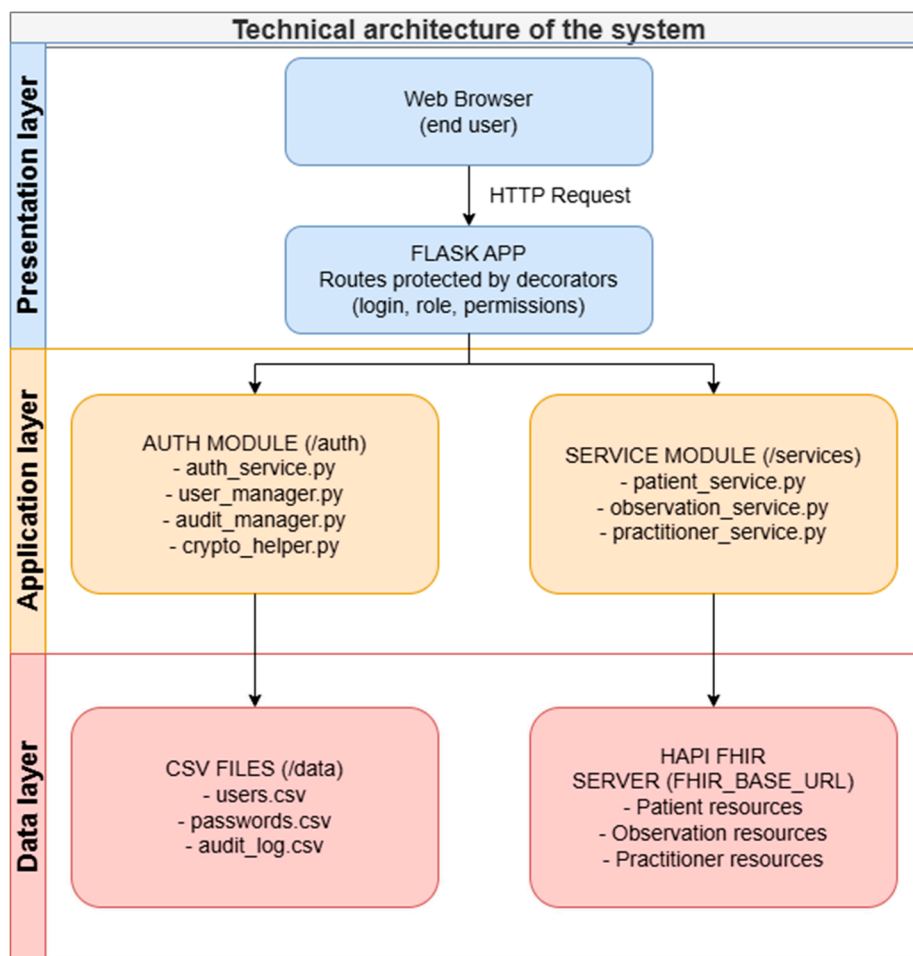


Figure 2. 3-layer architecture diagram.

The typical data flow in PIRE begins when an authenticated user sends an HTTP request (e.g., create an observation) through the web interface. Subsequently, the Flask application intercepts the request on a specific route, where the security decorator verifies that the user has the necessary role. If validation is successful, the controller calls the corresponding service, which executes the business logic where the CSV file is read and validated, the measurement type is automatically detected, the clinical interpretation is calculated, and the FHIR Observation resource is constructed according to the FHIR R4 standard. Following this, the constructed FHIR resource is transmitted via POST to the HAPI FHIR server endpoint, which persists it in its internal PostgreSQL database. Simultaneously, the authentication module records the action in the local file with user information, IP, timestamp, and transaction result, ensuring complete traceability. Importantly, this three-layer decoupled design allows changes to the user interface to not affect business logic, modifications to services to not require alterations to data access, and FHIR data persistence to be independent of local user management and auditing.

3.2.2. Functional Components and Modules

The technical implementation of PIRE structures its code into three specialized modules that operate in coordination under the control of the central Flask application. This modular organization allows each component to be responsible for a specific set of functionalities, thereby facilitating independent testing, incremental maintenance, and code reuse. Importantly, the modules interact through function calls within the application, avoiding direct coupling and allowing internal changes in one module to not affect others while maintaining their public interface. Table 2 below summarizes the three modules and their primary responsibilities.

Table 2. System modules and responsibilities.

Module	Components	Functional Responsibilities
auth/	Authentication, user manager, cryptography, security decorators, auditing	Identity validation; session creation and destruction; permission evaluation; secure encryption; credential generation; path protection; event auditing
services/	Patients, observations, professionals	Patient lifecycle (CRUD); CSV processing; automatic measurement detection; clinical validation; FHIR construction; remote synchronization
utils/	Logging, validators	Patient lifecycle (CRUD); CSV processing; automatic measurement detection; clinical validation; FHIR construction; remote synchronization

The Authentication Module (auth/) centralizes all functions related to security, access management, and system auditing. Specifically, this module is responsible for validating user identities against stored credentials, creating and destroying secure sessions, evaluating permissions based on assigned roles, protecting web paths through authorization mechanisms, encrypting and decrypting passwords to ensure they are never stored in plain text, automatically generating secure credentials (username and password) for new users, and comprehensively logging critical system events for complete traceability.

The module consists of five specialized files that work together. First, the authentication service (auth_service.py) provides the core logic allowing the application to verify who the user is and what actions they are authorized to perform. Second, the user manager (user_manager.py) handles the persistence of user data in CSV files, allowing you to create, query, update, and delete records securely and consistently. Third, the cryptographic assistant (crypto_helper.py) provides encryption functions that ensure passwords are never stored in plain text by using secure algorithms to protect credentials. Fourth, security decorators (decorators.py) define reusable mechanisms that protect individual application routes by automatically validating permissions before allowing access to functionality. Finally, the audit manager (audit_manager.py) logs critical authentication and user management events (successful logins, failed login attempts, logouts, user creation and deletion) with user information, IP address, timestamp, and result, creating a comprehensive audit log for regulatory compliance and subsequent forensic analysis.

The Services Module (services/) is responsible for orchestrating all operations involving clinical data, including managing the entire patient lifecycle, processing biomedical observations from CSV files generated by the Bioharness-3 device, generating FHIR R4 resources in accordance with international and Colombian regulatory standards, and ensuring resilient communication with the remote HAPI FHIR server.

A key feature of the module is its ability to automatically detect the type of measurement being analyzed. When a CSV file containing biomedical measurements is uploaded, the service automatically infers whether it corresponds to heart rate, respiratory rate, physical activity, or posture through statistical analysis of the data without requiring the user to manually specify the type. This capability significantly improves the user experience and reduces classification errors. Additionally, the module encapsulates all clinical validation logic, such as verifying that an identity document is unique in the system before creating a duplicate patient, validating that the patient's age is consistent with their document type, and calculating whether the measured values are within normal clinical ranges for automatic interpretation. The construction of FHIR resources occurs entirely within this module, ensuring that all data sent to the FHIR server complies with the defined structure, where each Patient resource includes specific Colombian extensions, each Observation includes standardized LOINC codes and accurate timestamps, and each Practitioner links correctly with system users.

The Utilities Module (utils/) provides reusable cross-cutting functionality that supports system integrity and traceability. Specifically, the logging subsystem configures a centralized logging system

that captures operational events at different severity levels, simultaneously writing to permanent files for later auditing and also to the console for debugging and real-time monitoring.

The validation subsystem implements a set of reusable validators that protect data integrity at multiple levels by verifying that text fields contain only expected characters, that email addresses have a valid structure with the correct domain, that phone numbers comply with regional patterns, that identity documents correspond to valid types in Colombia, and that dates have the correct format and are within the valid range. These validators are applied on both the client and server sides, ensuring that malformed data never leads to the creation of invalid FHIR resources that could cause errors on the HAPI FHIR server.

3.3. Authentication, Authorization, and Access Control System

The PIRE authentication and authorization system was designed as a layered security mechanism that ensures only identified and authorized users can access critical clinical management functions. Specifically, the solution combines credential-based authentication, role-based access control, application-level route protection, and comprehensive audit logging, aligning platform operation with recommended security practices for electronic health record systems.

3.3.1. Model of Roles Implemented

PIRE implements four main roles inspired by the organizational structure of a hospital environment: Super Admin, Doctor, Nurse, and Auditor, each with a clearly defined scope of authority over patient resources, observations, users, and audits. The Super Admin role acts as the overall system administrator, with the ability to perform comprehensive operations on patients and observations, manage users, view all user passwords, and access the full audit trail. Consequently, it is reserved for the initial configuration and overall administration of the platform.

The Doctor role represents the clinical professional with permissions to create, consult, and edit patients, record biomedical observations, and delete erroneous measurements, but without the authority to delete patients, administer users, or review audits, with the aim of separating administrative and care functions. In contrast, the Nurse role is geared toward nursing staff with an emphasis on operational records, where they can consult patients, create new observations, and review existing measurements, but cannot modify patient demographic data, delete observations, or manage users, which limits their impact on the structure of the clinical record. Finally, the Auditor role is designed for supervision and monitoring functions, with read-only access to patients and observations, and full access to the audit module, but without the ability to create, edit, or delete resources, ensuring independence between evaluation and operation. Table 3 provides detailed information on the permissions assigned to each role.

Table 3. Permission matrix by role.

Resource/Action	Super Admin	Doctor	Nurse	Auditor
Create patients	✓	✓	X	X
List patients	✓	✓	✓	✓
Consult patients	✓	✓	✓	✓
Edit patients	✓	✓	X	X
Delete patients	✓	X	X	X
Create observations	✓	✓	✓	X
Consult observations	✓	✓	✓	✓
See details of observations	✓	✓	✓	✓
Delete observations	✓	✓	X	X
Create user	✓	X	X	X
List user	✓	X	X	X
View user details	✓	X	X	X

View passwords	✓	X	X	X
Delete user	✓	X	X	X
View audit	✓	X	X	✓
View my profile	✓	✓	✓	✓

This role model embodies a principle of least privilege, whereby each actor only has the permissions strictly necessary for their usual tasks, reducing the attack surface and the risk of accidental or malicious modifications to sensitive clinical data. Moreover, the explicit separation between care, administrative, and supervisory roles facilitate the traceability of responsibilities in the event of incidents or internal audits.

3.3.2. Authentication Process and Session Management

Authentication in PIRE is based on individual credentials that are validated on the server before granting access to any protected functionality. As shown in Figure 3, the complete login and sign-in flow is orchestrated through multiple security checks. During system initialization, a startup script generates an initial Super Admin user whose credentials are stored in encrypted CSV files and documented in a restricted credentials file for configuration purposes only. When a user accesses the login page, the Flask application receives the credentials, verifies the user's existence in the users.csv table, and compares the password provided with the encrypted version in passwords.csv using a controlled decryption process.

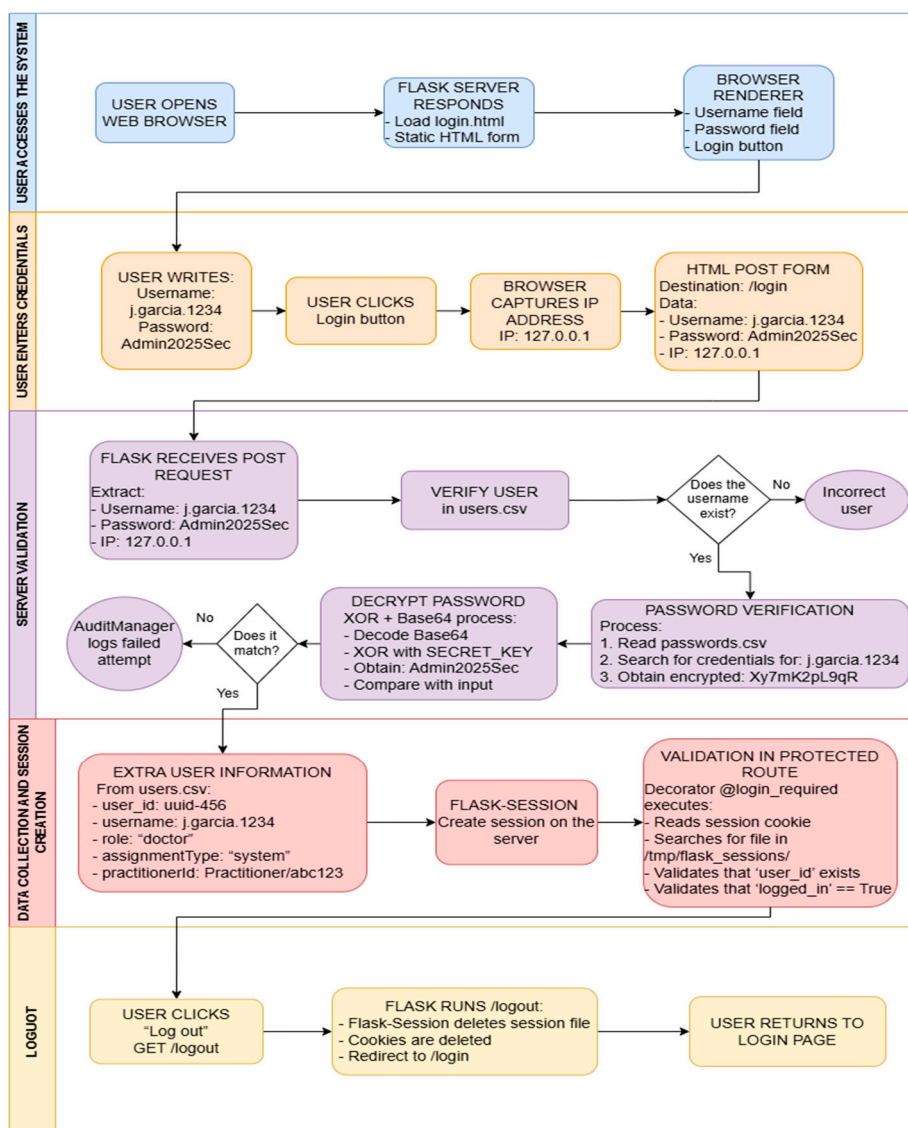


Figure 3. Login flow and session creation.

Upon successful authentication, the application creates a server session using Flask-Session, storing only internal identifiers such as the user ID, username, and associated role, without saving the password in the session or exposing sensitive information in cookies. To enhance security, sessions are stored in the server's file system, automatically invalidated when the application is stopped, and accompanied by HTTP headers that disable browser caching, reducing the risk of information exposure on shared computers. In each subsequent request, security decorators verify the presence of an active session and, if not, redirect to the authentication form, ensuring that no critical paths are accessible without login.

Notably, logging out explicitly invalidates the server session and records the event in the audit log, allowing you to identify each user's logins and logouts. Overall, this approach, based on server sessions and non-cache headers, complements the authentication model with additional measures against session reuse and screen exposure in previously authenticated browsers.

3.3.3. Authorization Mechanisms and Permission Matrix

On top of authentication, there is an authorization layer that evaluates, for each request, whether the user's role has explicit permission for the action they are attempting to perform. Specifically, this assessment is implemented based on the centralized permission matrix managed and shown in Table 3. The application associates each critical path with a logical resource and an action, so that the authorization layer can decide deterministically whether to grant or deny the operation. To implement this, role and permission security decorators integrate this logic into Flask routes, intercepting requests before business logic is executed. Specifically, the role-based decorator verifies that the session role is included in a set allowed for the path, while the permissions decorator consults the role permissions matrix to assess whether the role has rights over the resource and the requested action, returning a 403 error in case of unauthorized access. This declarative approach allows authorization logic to be kept concentrated in a readable permission structure, facilitating policy review and adjustment, and avoiding inconsistencies between different views that access the same type of resource.

3.3.4. Credential Security and Automatic User Generation

PIRE implements a specific scheme to protect passwords, combining an XOR operation on each character with a fixed secret key and subsequent Base64 encoding, so that passwords are never stored in plain text in the file system. During the encryption process, the plaintext password is transformed into an encoded representation that is stored in passwords.csv, while the users.csv file only retains user metadata such as the identifier, username, assigned role, creation date, and link to the corresponding Practitioner resource, thereby maintaining an explicit separation between identity and secrecy. Notably, password decryption is only performed at controlled moments, such as login validation or explicit viewing by the Super Admin, applying the reverse process of Base64 decoding and XOR operation with the same secret key.

User generation is supported by an automated workflow designed to reduce human error and standardize credentials. When the Super Admin creates a new user through the administration interface, the platform requires the completion of a form with the user's personal data, including first and last names, type and number of identification document, contact details, specialty, and FHIR role. After validating that the identity document is unique and that the role and specialty are consistent, the system automatically creates a Practitioner resource on the FHIR server, generates a username following a deterministic pattern, and produces a secure random password between 10 and 15 characters that must include uppercase letters, lowercase letters, numbers, and symbols.

Subsequently, new user information is stored in users.csv and their encrypted password in passwords.csv, with the operation being recorded in the audit log along with the identifier of the associated FHIR professional. To enhance security, the generated credentials are displayed only once

in the Super Admin interface, who must transmit them to the professional through a secure channel, preventing the password from being automatically displayed again to reduce the risk of inadvertent exposure. Additionally, the platform offers a “My Profile” view accessible to all authenticated users, allowing them to review their basic data, assigned role, and credentials, thereby reinforcing the transparency of the access model.

3.3.5. Auditing Authentication and Access Events

The PIRE audit system is the final layer of the security model, providing complete traceability of authentication events, access, and management of clinical and administrative resources. Specifically, each login attempt, whether successful or unsuccessful, generates a record in the `audit_log.csv` file with information about the user, the action performed, the IP address, the timestamp, and an indicator of success or failure, allowing patterns of failed attempts or accesses from unusual locations to be identified. Similarly, all sensitive user and professional management operations, such as the creation or removal of Practitioner users and resources, are recorded with the action detail, the type of resource concerned and its identifier.

The audit interface presents the last one hundred records arranged chronologically, thereby allowing for quick identification of recent actions, and highlights different types of events through visual indicators, facilitating anomaly detection, incident investigation, and compliance reporting, making auditing an active component of the platform’s governance model. Taken together, the combination of robust authentication, role-based and permission-based authorization, credential protection, and detailed auditing creates a comprehensive security environment consistent with the needs of an interoperable clinical management system.

3.4. Patient Management and Resource Generation FHIR Patient

Patient management in PIRE was implemented as a complete lifecycle flow on the FHIR R4 Patient resource, aligned with the FHIR Core CO Implementation Guide. Specifically, the web application offers structured forms for patient registration, consultation, updating, and deletion, while the HAPI FHIR server acts as a central clinical repository, storing each patient as a Patient resource enriched with extensions specific to the Colombian context. Critically, the integration between the Flask interface and the FHIR server ensures that all demographic and identification information subsequently used by biomedical observations remains consistent, versioned, and accessible for clinical and secondary analysis purposes.

3.4.1. Processes for Creating, Reading, Updating, and Deleting Patients (CRUD)

Registering a new patient is done through a multi-step web form that guides the user through capturing the minimum data required by the FHIR Patient model, including official identifiers, demographics, contact information, and emergency contact details. During this process, the system validates that the identity document does not already exist on the FHIR server, applies format validators to names, emails and telephones, and verifies the consistency between age and document type, avoiding the creation of duplicate or inconsistent records. Once the validations are passed, the application builds a Patient resource and sends it via a POST operation to the `/Patient` endpoint of the HAPI FHIR server, which returns the patient’s unique identifier for reference for future clinical operations.

The consultation and listing of patients is supported through searches by identity document and paginated lists that obtain information from the FHIR server using standard search parameters. To facilitate clinical use, the interface displays full name, document type and number, and FHIR identifier, allowing the user, with permissions on patient data, to navigate to the detailed patient record, where complete demographic data is displayed and linked biomedical observations are accessed. With respect to modifications, update operations are restricted to roles with clinical permissions, which can modify specific contact fields or demographic data. These changes are

reflected through update operations on the Patient resource in the FHIR server, maintaining a single master record per patient. Notably, deleting patients is reserved for the Super Admin role and is implemented as a cascading process that first removes the observations associated with the patient and subsequently removes the Patient resource from the FHIR server, reducing the risk of leaving orphaned clinical data. As illustrated in Figure 4, the CRUD operational flow of the Patient resource demonstrates the complete lifecycle.

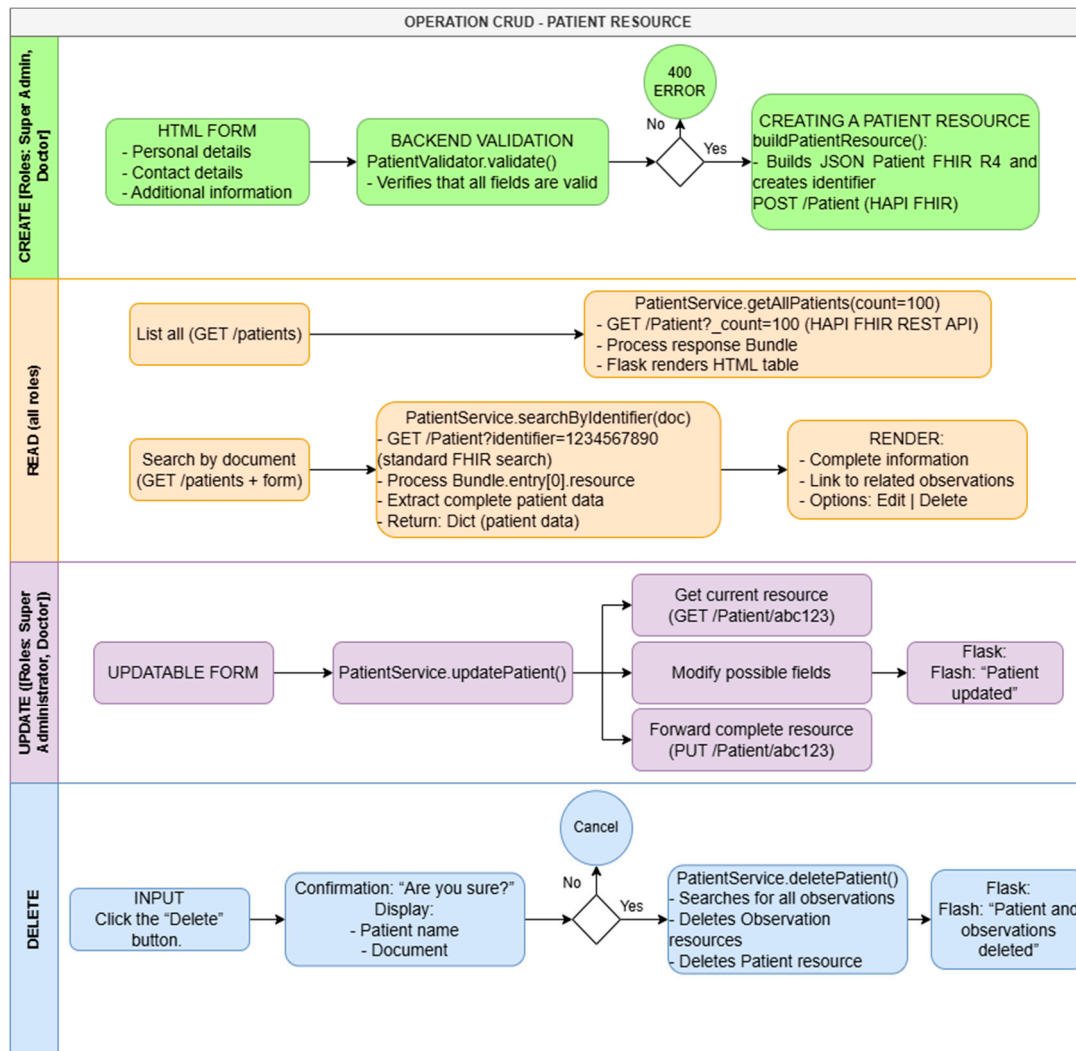


Figure 4. CRUD operational flow of Patient resource.

3.4.2. FHIR Patient Resource Structure

The Patient resource generated by the platform follows the FHIR R4 model, incorporating the essential elements for unique identification and clinical use in the Colombian context. Specifically, each resource includes an official identifier in the identifier element, which combines the document type and the patient's identification number, allowing deterministic searches from the application forms. To preserve regional naming conventions, the patient's name is represented with explicit separation between first and last names, and with support for registering the second maternal surname through a specific extension, preserving the typical nominal structure of the region. Additionally, the Patient resource includes coding of gender, date of birth, contact information and address structured by lines, city and department, which allows its reuse in both care flows and subsequent geospatial analysis. Furthermore, marital status is also stored and, when available, occupation information via extensions, as well as an emergency contact with name, relationship and contact details, which facilitates continuity of care and communication in critical events. Collectively,

this standardized structure ensures that demographic data is interoperable and that Patient resources can be consumed by other FHIR-compliant systems without additional transformations.

3.5. Biomedical Observation Taking and Resource Generation FHIR Observation

PIRE implements a specialized workflow to record biomedical observations from CSV files generated by the Bioharness-3 device, transforming time series of physiological signals into FHIR Observation resources structured according to the R4 standard. The process combines guided web forms, a robust CSV file processing engine, and clinical business logic that ensures data validation, automatic detection of measurement type, and stratification of values according to preconfigured interpretive ranges. In this manner, each set of measurements of heart rate, respiration, physical activity or posture is represented as an interoperable observation, linked to a previously registered patient and available for clinical consultation and further analysis.

3.5.1. CSV File Upload and Data Validation Process

The registration of a new observation is done through a two-step flow that begins with the selection of the patient and the measurement context, followed by the uploading of the CSV file with the biomedical measurements. **In the first step**, the authenticated user selects the target patient from their identity document, validates their existence on the FHIR server, and specifies contextual information such as the device used (Bioharness-3) and the expected signal type, temporarily storing these parameters in the application session. **In the second step**, the interface allows you to upload the CSV file exported from Bioharness-3, which is generated from the Software Development Kit (SDK) provided by the company Zephyr called Omnisense, and fill in additional clinical metadata, after which the observations module runs a series of validations on the received file before building the FHIR resource.

The CSV processing system was designed to tolerate frequent variations in encoding, separators, and time formats typical of real-world environments. Specifically, for each file, the service attempts to decode the content using different encoding schemes and test multiple separators until it identifies a valid combination that allows the columns to be interpreted correctly. The expected format requires a first header line, which is discarded, a first column with timestamps, and a second column with numerical values corresponding to the recorded physiological signal. To ensure temporal integrity, timestamps are normalized by applying a set of predefined patterns, ensuring that all points in the series are converted to valid time instants before constructing the observation.

3.5.2. Automatic Detection of Measurement Type

A central feature of the observations module is the ability to automatically detect the type of measurement contained in the CSV file through statistical analysis of the numerical values. The service calculates basic metrics such as range of values, average and presence of decimals and, from these parameters, classifies the signal into one of four categories: heart rate, respiratory rate, physical activity, posture. This approach enables to identify, for example, continuous value series around tens or hundreds such as heart rate, lower and regular values such as respiratory rate, numerical levels typical of physical activity scales, and discrete codes that represent changes in posture.

Notably, automatic detection reduces the user's dependence on correctly declaring the measurement type and decreases the likelihood of classifying a signal under an inappropriate clinical code. In cases where the statistical detection does not match the type declared by the user in the first step, the system can mark the observation as inconsistent or warn the user, avoiding the creation of FHIR resources with ambiguous semantics. Importantly, this logic is integrated before the construction of the Observation resource, so that observations are only generated when there is consistency between the data in the file and the type of measurement identified.

3.5.3. Construction of the FHIR Observation Resource

Once the CSV file has been validated and the type of measurement determined, the platform builds a FHIR Observation resource that encapsulates the time series of measurements using the SampledData structure. Specifically, the code element of the observation is encoded with LOINC terminologies specific to each type of signal, using codes such as 8867-4 for heart rate, 9279-1 for respiratory rate, 41950-7 for physical activity and 8361-8 for posture, which ensures that the semantics of each observation are recognizable by other interoperable systems. The resource links to the corresponding patient through the subject element, referencing the FHIR Patient identifier and maintaining the explicit relationship between the measurements and the patient's master record.

The series of values is represented by the SampleData type in the value-SampledData element, which includes parameters such as the initial time, the sampling interval, the unit of measurement, and the sequence of space-separated values, derived from the second column of the CSV file. Additionally, the resource also includes temporary metadata where the date of creation of the observation in the system is recorded, which allows observations to be grouped by upload date in the query interface. When additional information about the device or measurement context is available, this data is incorporated into complementary elements or extensions, strengthening the traceability of observations to the physical origin of the signals.

3.5.4. Clinical Interpretation Calculation and Use of Ranges

In addition to storing the raw signal values, the platform calculates a summarized clinical interpretation for each observation, relying on a configuration file of normal ranges. During CSV file processing, the service determines a representative value of the series, typically the average of the measurements, and compares it with reference ranges specific to each type of signal and LOINC code. Based on this comparison, it assigns a categorical classification, providing an immediate summary of the patient's condition in the context of that measurement. This interpretation is stored as part of the Observation resource, allowing both the user interface and other FHIR-consuming systems to exploit the information without the need to permanently recalculate clinical ranges. In consultation views, the platform uses this classification to highlight out-of-range observations and make it easier for the professional to quickly identify episodes of tachycardia, bradycardia, respiratory disturbances, or unusual activity and posture patterns. Notably, separating raw data from interpretation criteria into a separate, configured file also allows for adjusting ranges without modifying the system logic, thereby facilitating adaptation to different clinical protocols or target populations.

3.5.5. Error Handling

Error handling during CSV file processing was implemented as a cross-cutting component of the observations module, aimed at preserving data integrity and providing clear feedback to the user in the face of loading failures. If none of the configured encoding schemes manage to decode the file, or if the separators do not allow the expected column structure to be correctly reconstructed, the platform catches the exception, records the failed attempt in the application logs, and displays an error message indicating that the file could not be processed, suggesting verification of the encoding and format of the source file. Similarly, when automatic detection of the measurement type does not produce a reliable classification or is inconsistent with the type declared by the user in the first step of the flow, the system marks the observation as invalid, stops the construction of the FHIR Observation resource and requests a review of the data, preventing measurements with ambiguous semantics from persisting.

In cases where communication errors occur with the HAPI FHIR server during the POST operation of the observation, the service implements automatic retries with increasing wait times and records each attempt in the log to differentiate transient connectivity problems from structural failures in the FHIR server. To mitigate security risks, the uploaded CSV files are temporarily stored in the application's working directory and are securely deleted once the observation has been

successfully created on the FHIR server or when the process is aborted due to unrecoverable errors, reducing the risk of file accumulation and exposure of sensitive data in the file system. As shown in Figure 5, the complete process for ingesting medical observations demonstrates the integrated workflow from upload to FHIR resource creation.

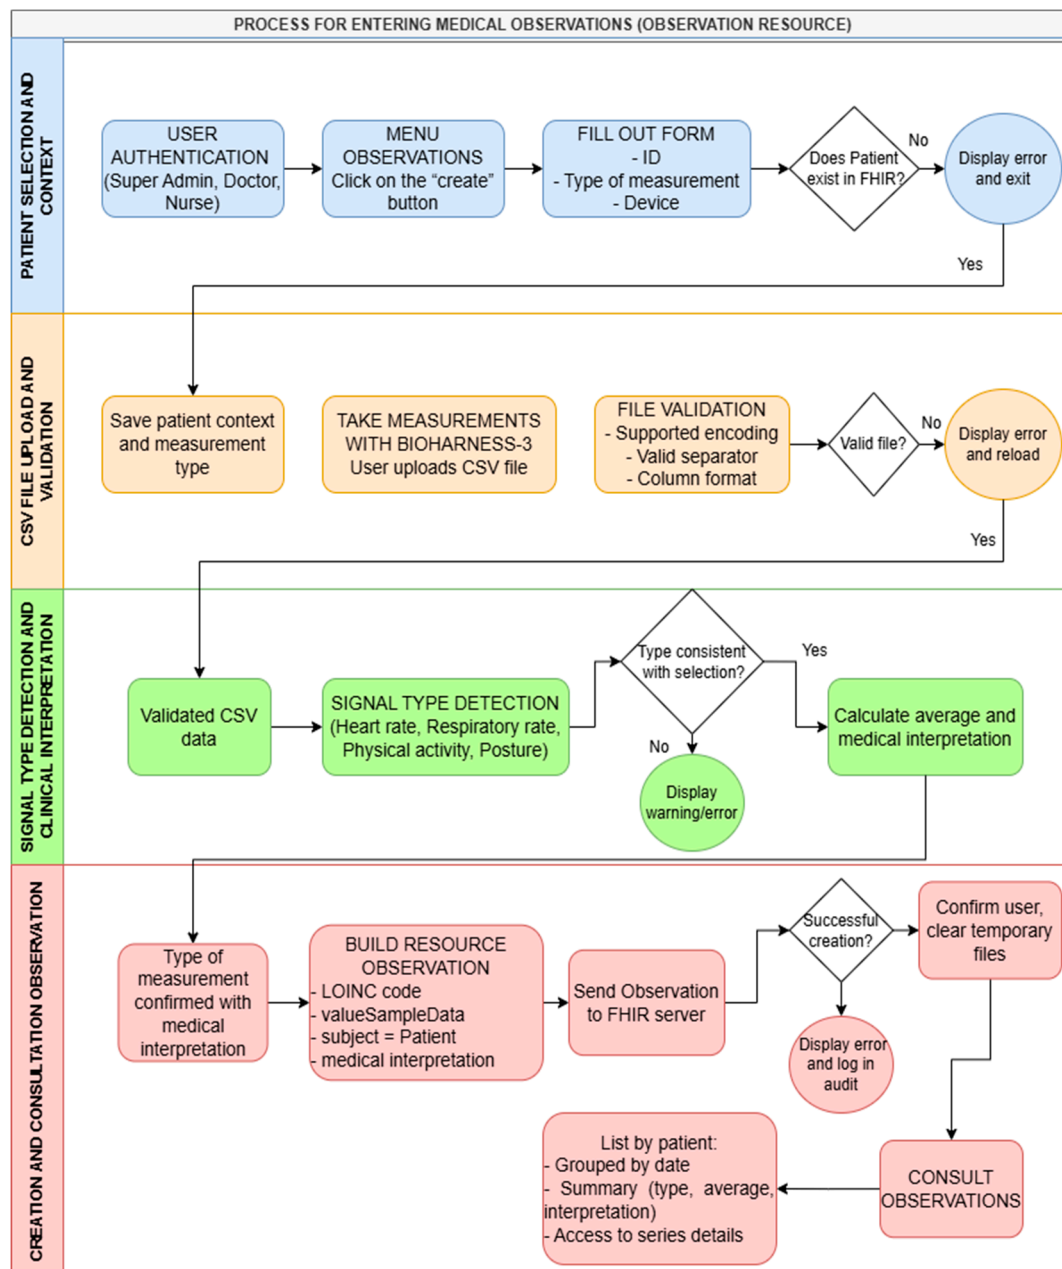


Figure 5. Medical observation intake process – Observation resource creation.

3.6. FHIR Practitioner Health Professional and Resource Management

The management of health professionals in PIRE was modeled using the FHIR R4 Practitioner resource as a representation of the clinical professional, while the internal user system provides the access credentials and associated security roles. **This separation** allows for a clear distinction between clinical identity and access identity, while ensuring an explicit link between the two by storing the Practitioner identifier in the user record.

3.6.1. Creation of the Practitioner Resource

The creation of a new healthcare professional in the system is done in an integrated way with the user flow managed by the Super Admin role. When a new clinical user is registered through the user management form, the platform requests personal and professional data such as names, surnames, type and number of document, contact information, city and department, as well as the specialty or position they hold in the institution. After validating this data, the backend invokes the `practitioner_service` to build and send a Practitioner resource to the HAPI FHIR server, using a POST operation on the `/Practitioner` endpoint.

The generated Practitioner resource includes an identifier based on the professional's document, the full name structured in first and last names, and when available, attributes such as gender, date of birth, entered phone numbers and emails, in addition to the contact address. The specialty or position is represented in the Practitioner's professional qualification elements, allowing other systems to know the clinical profile of the professional who creates or validates observations and manages patients. Once the resource is created, the FHIR server returns the unique identifier that is stored along with the user in the `users.csv` file and is subsequently used to associate clinical actions with the corresponding professional, for example, as creator in biomedical observations.

3.6.2. Integration with the User System

The integration between the Practitioner resources and the user system is based on a direct link through the `practitioner_id` field stored in the `users.csv` file, which stores, for each user, the identifier of the associated Practitioner resource in the FHIR service. In this manner, each user account with a role (Doctor, Nurse, or Auditor) maintains a stable reference to the healthcare professional it represents, allowing the platform's functionalities to retrieve and display clinical information about the professional when necessary. In the "My Profile" view, the system queries the associated Practitioner to present the user with their personal data, document, specialty, and FHIR identifier, reinforcing the transparency of the link between access identity and clinical identity.

This coupling is also used to ensure consistency in the clinical actions recorded on the FHIR server, since observations and other operations can reference the responsible professional through the Practitioner identifier stored alongside the authenticated user. When a user is deleted from the administration interface, the platform executes a cascading process that includes the deletion of the corresponding Practitioner resource on the FHIR server, as well as the removal of its records in `users.csv` and `passwords.csv`, logging the entire operation in the audit log to preserve traceability. Consequently, this strategy ensures that no orphan Practitioner resources remain associated with non-existent user accounts and keeps the clinical FHIR repository, and the internal user management system aligned.

3.7. System Auditing and Logging

PIRE incorporates an auditing and logging subsystem designed to record, in a structured manner, the critical actions performed by users and the relevant events of the application. This subsystem combines a functional log oriented toward compliance and clinical traceability, stored in the `audit_log.csv` file, with a technical log in application logs that allows diagnosing errors, monitoring system behavior, and supporting maintenance tasks.

3.7.1. Logging of Critical Actions

The auditing module was implemented through the `AuditManager` class of the `/auth` component, which exposes operations to record and query audit events from a persistent CSV file. Each critical action, such as successful or failed logins, user creation and deletion, and login and logout, generates an entry in `audit_logs.csv` with structured information about the context of the operation. The audit file format includes fields such as record identifier, user identifier, username, action type, type of resource involved, resource identifier, additional details, source IP address,

timestamp, and a success or failure indicator, which allows accurately reconstructing the history of interactions with the system.

High-level functions and clinical service operations invoke the Audit Manager at key points in the flow, ensuring that each authentication attempt or modification of Practitioner resources is recorded. The audit interface, accessible only to the Super Admin and Auditor roles, queries the records through dedicated methods that allow retrieving the last one hundred events, filtering by user or by action type, and displaying the information in a table ordered chronologically. This view provides the responsible team with an immediate tool to review recent system activity, identify patterns of failed access attempts, and support formal auditing or incident investigation processes.

3.7.2. Transaction Traceability and Application Logging

Complementary to the functional audit file, the platform uses an application logging system configured through the `utils.logger` module, which records general operation events, errors, and security events in an `app.log` file located in the logs directory. This logger centralizes messages at different severity levels, including information on normal actions, warnings for anomalous situations, and execution errors with context, thereby facilitating problem diagnosis during development, testing, and operation. Among others, events such as connection failures with the HAPI FHIR server, errors in CSV file processing, unhandled exceptions, and relevant security events, such as repeated access attempts with incorrect credentials, are recorded.

The combination of functional auditing and technical logging provides traceability at both the clinical and infrastructure levels. Notably, while `audit_log.csv` is oriented toward answering questions about who did what, when, and on which action, the `app.log` file makes it possible to understand how the application behaved internally when executing those actions, what errors occurred, and in what context they appeared. Collectively, this duality facilitates compliance with governance and oversight requirements in clinical management systems, while also providing the technical team with concrete tools for problem resolution, early incident detection, and continuous improvement of the platform.

4. Results

The PIRE platform was developed as an application structured into nine main interfaces that allow users to interact with the different components of the interoperability system. These interfaces include comprehensive patient management, the creation and visualization of clinical observations derived from wearable devices, the administration of system users and roles, and an audit module that records all critical operations performed by the application. Through these interfaces, the platform demonstrates its operational capacity to transform clinical data from diverse sources into structured information in accordance with international standards and Colombian regulations. Below, the operation and concrete results of each of these interfaces during the technical validation phase are described. It is important to emphasize that, in order to protect the anonymity of registered patients and users, example names and credentials were used to illustrate the interfaces.

4.1. Patient Management

Patient management constitutes the fundamental entry point for the registration of subjects of care in PIRE. Specifically, this section allows the creation of new patient records, the visualization of the complete list of patients registered in the system, the search for specific patients by ID, and the deletion of records when necessary. The design was conceived to capture the essential demographic data required by both Colombian clinical interoperability regulations and the FHIR R4 standards, ensuring that each patient entered into the system automatically generates a FHIR Patient resource in accordance with the PatientCO profile of the FHIR Core Colombia Implementation Guide.

4.1.1. Patient Creation Interface

The PIRE patient creation interface presents a form structured into three clearly differentiated sections, designed to guide the user in filling in the minimum required information and the complementary data of the subject of care. In the first section, called Mandatory Data, the essential fields for patient identification are captured: first names, last names, gender, identification number, document type, and date of birth. With this combination of attributes, the system has sufficient information to construct a Patient resource in accordance with the guidelines of the FHIR Core Co Implementation Guide, respecting Colombian identification conventions and adequately representing the person's identity within the clinical interoperability ecosystem. As shown in Figure 6, the mandatory data section of the form captures these critical identifiers.

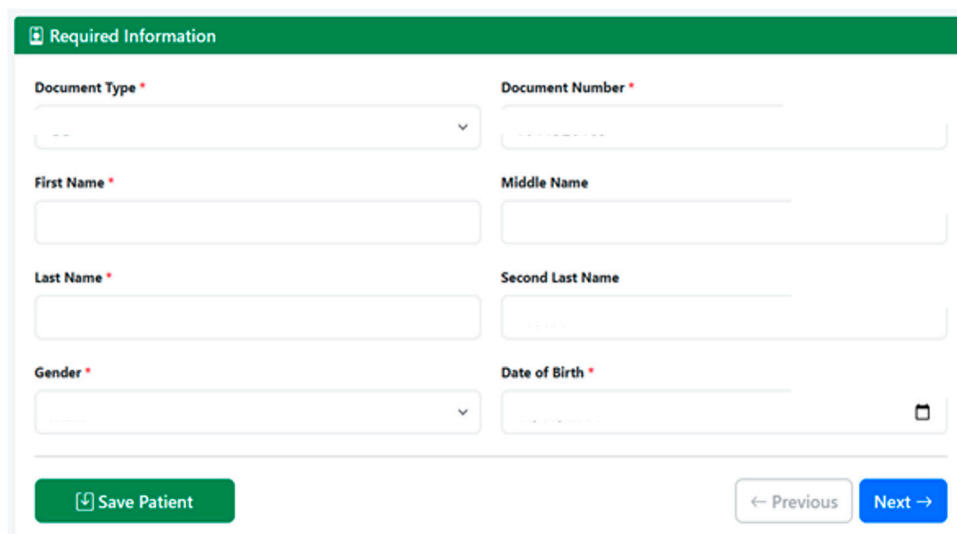


Figure 6. Mandatory data section in the patient creation form.

The second section of the form corresponds to the Contact Data, where the phone number, email address, residential address, city, and state or province are recorded. These fields make it possible to enrich the Patient resource with telecommunications and location information that is fundamental for both care processes and for notifications, follow-up, and territorial analysis. As illustrated in Figure 7, the contact data section collects these communication parameters.

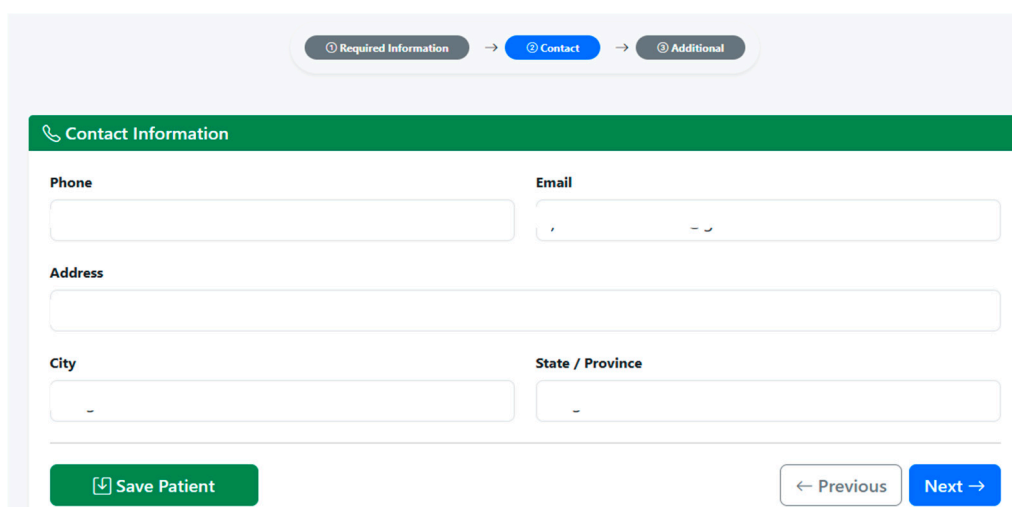


Figure 7. Contact data section of the patient creation form.

Finally, the third section, called Additional Information, collects complementary attributes such as marital status, patient occupation, emergency contact, and emergency phone number. This information not only provides relevant clinical and social context, but also facilitates support in critical situations. As demonstrated in Figure 8, the additional data section completes the demographic profile. Once these three sections have been completed and the consistency of the fields has been validated, the system proceeds to generate and store the corresponding FHIR Patient resource on the interoperability server.

Figure 8. Additional data section of the patient creation form.

4.1.2. Automatic Generation of the FHIR Patient Resource

Each time a patient is successfully created in the patient creation interface, the system automatically builds a FHIR Patient resource that incorporates not only the basic demographic information but also the extensions and specific elements required by the PatientCO profile. As illustrated in Figure 9, the complete structure of a Patient resource generated by the platform demonstrates the semantic richness and compliance with Colombian regulatory standards.

```

{
  "resourceType": "Patient",
  "id": "b7813bc5-5d2d-47c4-9204-917fb42cf44f",
  "meta": {
    "profile": [
      "https://www.minsalud.gov.co/ihc/fhir/StructureDefinition/PacienteCO"
    ]
  },
  "identifier": [
    {
      "use": "official",
      "type": {
        "extension": [
          {
            "url": "https://www.minsalud.gov.co/ihc/fhir/StructureDefinition/co-pais",
            "valueCodeableConcept": {
              "coding": [
                {
                  "system": "https://www.minsalud.gov.co/ihc/fhir/CodeSystem/PaisesCS",
                  "code": "CO",
                  "display": "CO"
                }
              ]
            }
          }
        ]
      },
      "coding": [
        {
          "system": "https://www.minsalud.gov.co/ihc/fhir/CodeSystem/IdentificadoresPersonaCS",
          "code": "CC",
          "display": "Cédula de ciudadanía"
        }
      ]
    }
  ],
}

```

(a)

```

    "value": "1011320189"
  ],
  "active": true,
  "name": [
    {
      "use": "official",
      "family": "Barbosa",
      "given": [
        "Juan"
      ]
    }
  ],
  "gender": "male",
  "birthDate": "2004-02-11",
  "extension": [

```

(b)

Figure 9. Complete FHIR Patient Resource: (a) first part of the FHIR Patient resource, (b) second part of the FHIR Patient resource.

This resource is automatically sent to the HAPI FHIR server and registers the user on the platform through an HTTP POST request performed internally by PIRE. The structure includes all the data that are filled in the patient creation form, where the mandatory data are the data that must be present in the Patient resource without exception.

4.1.3. Patient Listing and Visualization Interface

The patient listing interface presents a table that displays the complete set of patient records created in the system. During the technical validation of PIRE, the interface displayed all registered patients organized in columns that include the patient's full name and identification document number. The system allows users with the permitted role to perform quick searches using the patient's document number, facilitating the location of specific patients without the need to browse the entire list. As shown in Figure 10, the patient listing and validation interface provides consolidated access to the patient database.

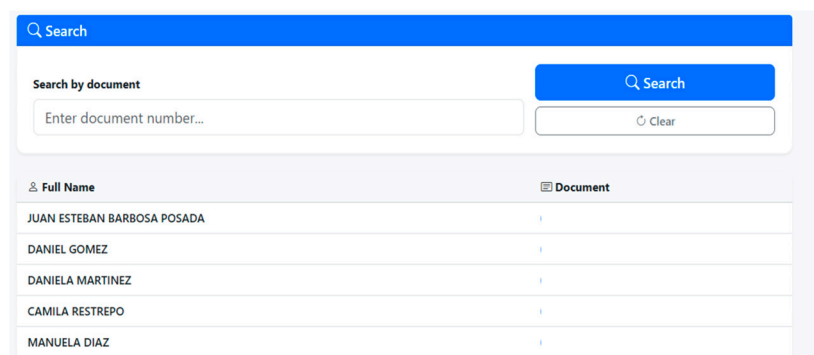


Figure 10. Patient listing and visualization interface.

4.1.4. Patient Search Interface

The implemented search interface made it possible to perform dynamic queries on the patient database using the patient's identification document number, thereby allowing the visualization of all the patient's recorded information. This functionality was validated during the technical validation phase, making it possible to verify the system's ability to retrieve specific records efficiently. When a user performs a patient search, they enter the identification document number and all the information that was previously filled in the patient creation interface is displayed. As illustrated in Figures 11 and 12, the patient search interface demonstrates successful retrieval of patient demographic data.

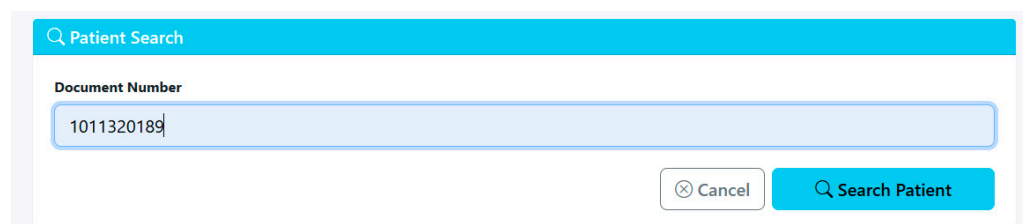


Figure 11. Patient search by identification document number in the patient search interface.

Patient Information

PATIENT ID: 1

DOCUMENT: [REDACTED]

Personal Information

First Name: [REDACTED] Middle Name: [REDACTED] Last Name: [REDACTED] Second Last Name: [REDACTED]

Gender: [REDACTED] Date of Birth: [REDACTED] Marital Status: [REDACTED]

Contact

Phone: [REDACTED] Email: [REDACTED]

Address

Address: [REDACTED] City: Bogotá State: Bogotá

Additional Information

Occupation: Engineer Emergency Contact: [REDACTED]

Figure 12. Information of the queried patient in the patient search interface.

4.1.5. Patient Deletion Interface

The patient management section included the functionality for deleting records, accessible only to users with the permitted role. When a deletion operation is executed, the system is responsible for completely removing all patient data, including observations that are related to the patient through their identification document number. As shown in Figure 13, the patient deletion interface enforces role-based access controls to protect data integrity.

Warning

This action will permanently delete the patient from the FHIR system. This operation cannot be undone.

Delete Patient

Patient Document Number

26789345

Enter the document number of the patient you want to delete

Cancel Delete Patient

Figure 13. Patient deletion interface.

4.2. Biomedical Observations Management

The management of observations of the PIRE platform combines guided multi-step web forms, integrated data validation, automatic detection of measurement type, and stratification of values according to preconfigured interpretive ranges. In this manner, each set of heart rate, respiratory rate, physical activity, or posture measurements is represented as an interoperable observation, linked to a previously registered patient, associated with the biomedical device that originated the measurement, and available for clinical consultation, comparative analysis, and data source traceability.

4.2.1. Observation Creation Interface

The observation creation interface in PIRE guides the user through a flow of four sequential stages, each capturing specific information necessary to build a complete and semantically correct Observation resource. In the first stage, called observation data, the user enters three fundamental data points: the identification number of the patient to whom the observation will be associated, the biomedical device used for the measurement (specifically Zephyr Bioharness-3 in the technical validation of PIRE), and the type of measurement to be entered, selecting between heart rate, respiratory rate, physical activity, or posture. This information is temporarily stored in the user session to guide the subsequent steps of the flow. As illustrated in Figure 14, the first stage of the observation creation interface establishes the measurement context.

The screenshot shows a form titled "Observation Data" with three numbered steps. Step 1 is "Enter the patient's document number *", with a text input field containing "1011320189" and a tooltip that says "If the patient does not exist, the observation cannot be created". Step 2 is "Select the observation device *", with a dropdown menu showing "Zephyr BioHarness-3" and a tooltip that says "Device used to capture the measurements". Step 3 is "Select the measurement type *", with a dropdown menu showing "Heart Rate" and a tooltip that says "Measurement type contained in the CSV file".

Figure 14. Observation data stage of the observation creation interface.

In the second stage, called information validation, the system queries the FHIR server to verify that the identified patient exists in the clinical repository and displays the patient's demographic information (identification number and full name) together with a confirmation of the biomedical device data and the selected measurement type. This validation ensures that, before proceeding with the ingestion of clinical data, there is an explicit correspondence between the observation and the subject of care, avoiding the creation of measurements associated with non-existent or incorrectly identified patients. The user must confirm this correspondence before continuing, strengthening traceability and accountability in the ingestion of clinical data. As shown in Figure 15, the information validation stage prevents orphaned observations.

The screenshot shows a form titled "Validated Information" with a green header. Below the header is a table with three columns: "ID", "FULL NAME", and "DEVICE". The "DEVICE" column contains "Zephyr BioHarness-3". Below the table, the "MEASUREMENT TYPE" is displayed as "Heart Rate".

Figure 15. Information validation stage of the observation creation interface.

In the third stage, called CSV information, the user completes complementary fields that contextualize the observation within the clinical process. These fields include the measurement method used (continuous monitoring, ambulatory monitoring, or high physical activity monitoring), the body site where the measurement was performed, and additional notes that the clinical professional considers relevant for the subsequent interpretation of the observation, such as patient conditions, interferences during the measurement, or specific clinical context. This information is captured as free text and will later be stored as part of the structure of the Observation resource, facilitating the reconstruction of the measurement context during consultations. As illustrated in Figure 15, the CSV information stage enriches observations with clinical context.

The screenshot shows a form titled "CSV Information" with a light blue header. Below the header, there are three main sections: "Measurement Method", "Measurement Site", and "Additional Notes". Each section has a dropdown menu or a text input field. The "Measurement Method" dropdown is set to "Ambulatory monitoring". The "Measurement Site" dropdown is set to "Chest". The "Additional Notes" section has a text input field with the placeholder text "Additional observations about the measurement...".

CSV Information: The CSV file will provide measurement values and timestamps. Complete the following fields to enrich the observation.

Measurement Method *

Ambulatory monitoring

Method used to capture the measurement

Measurement Site *

Chest

Location of the BioHarness-3 sensor on the body

Additional Notes *

Additional observations about the measurement...

Any relevant information about the measurement context

Figure 15. CSV information stage of the user creation interface.

In the fourth stage, called measurement CSV file, the user selects and uploads the file exported from the Zephyr Omnisense SDK that contains the time series of measurements. Once the file is uploaded, the platform performs validations on the CSV content, processes the information, and, if all validations are successful, builds and persists the Observation resource on the HAPI FHIR server. When this operation is completed, the application displays a confirmation message indicating "Observation created successfully" and shows the unique identifier generated by the FHIR server, providing the user with an explicit reference to the created resource that can be used for subsequent follow-up. Figure 16 illustrates the measurement CSV file upload stage.

The screenshot shows a form titled "Upload CSV File" with a yellow header. Below the header, there is a "Measurement CSV File" section with a file selection button and a file name "Tasa del corazon.csv". Below the file name, there is a green box indicating the selected file. Below the green box, there is a yellow box with "CSV Requirements" listed as bullet points.

Upload CSV File

Measurement CSV File *

Seleccionar archivo Tasa del corazon.csv

The file must contain columns X (timestamp) and Y (value)

Selected file: Tasa del corazon.csv (5,01 KB)

CSV Requirements:

- ✓ Format: CSV with tab or comma separator
- ✓ First row: Metadata (automatically ignored)
- ✓ Column X: Measurement timestamps
- ✓ Column Y: Numeric measurement values
- ✓ Data type must match the selected type

Figure 16. Measurement CSV file upload stage of the user creation interface.

4.2.2. Automatic Generation of the FHIR Observation Resource

Each time an observation is successfully created in PIRE, the system automatically builds an FHIR Observation resource that encapsulates the time series of biomedical measurements using the SampleData structure according to R4 specifications. Additionally, each type of signal is encoded with specific LOINC terminologies, using the specific LOINC codes for each measurement, ensuring that the semantics of each observation are recognizable by other interoperable systems. The resource links the corresponding patient by referencing the FHIR Patient identifier and maintaining the explicit relationship between the measurements and the patient master record. The data value series includes parameters such as the initial instant of the series, the sampling interval in milliseconds, the specific unit of measure for each type of signal, and the space-separated sequence of values derived from the CSV file. Finally, the resource contains all the additional information filled in the form and the diagnosis performed according to the obtained measurements. As illustrated in Figure 17, the Observation resource generated by the PIRE platform demonstrates complete semantic compliance.

```

(a)
{
  "resourceType": "Observation",
  "id": "a8c4df44-765c-4763-8f69-932fcaa77856",
  "meta": {
    "profile": [
      "https://loinc.org/41950-7/"
    ]
  },
  "status": "final",
  "category": [
    {
      "coding": [
        {
          "system": "http://terminology.hl7.org/CodeSystem/observation-category",
          "code": "vital-signs",
          "display": "Vital Signs"
        }
      ]
    }
  ],
  "code": {
    "coding": [
      {
        "system": "http://loinc.org",
        "code": "41950-7",
        "display": "Physical activity level"
      }
    ]
  },
  "subject": {
    "reference": "Patient/b7813bc5-5d2d-47c4-9284-917fb42cf44f",
    "display": "Juan Barbosa"
  },
  "effectivePeriod": {
    "start": "2025-09-22T16:57:27.499Z",
    "end": "2025-09-22T17:00:40.139Z"
  },
  "issued": "2025-12-01T16:49:27.909Z",
  "device": {
    "display": "Zephyr Bioharness-3"
  },
  "valueSampledData": [
    {
      "origin": {
        "value": 0,
        "unit": "MET",
        "system": "http://unitsofmeasure.org",
        "code": "MET"
      },
      "interval": 2501.818,
      "intervalUnit": "ms",
      "factor": 1.0,
      "lowerLimit": 0.01,
      "upperLimit": 0.36,
      "dimensions": 1,
      "data": "0.36 0.08 0.04 0.12 0.09 0.27 0.17 0.08 0.03 0.02 0.18 0.19 0.09 0.08 0.24 0.0"
    }
  ],
  "interpretation": [
    {
      "coding": [
        {
          "system": "http://terminology.hl7.org/CodeSystem/v3-ObservationInterpretation",
          "code": "L",
          "display": "Low"
        }
      ],
      "text": "Nivel de actividad muy bajo (promedio: 0.07 MET). Se recomienda incrementar"
    }
  ]
}

```

Figure 17. Complete FHIR Observation resource: (a) first part of the FHIR Observation resource, (b) second part of the FHIR Observation resource.

4.2.3. Observation Listing and Search Interface

The observation listing interface allows authorized users to query all measurements associated with a specific patient by searching by the patient's identification number. When a user enters the identification number in a search field, the system queries the FHIR server to retrieve all Observation resources linked to that patient, and presents a table with the consolidated information of each observation. The table includes columns such as the observation creation date, the measurement time, the type of observation, the interpretation of the observation, the average value of the measurement, and the option to view the observation details. This tabular presentation provides an executive view of the measurements performed on the patient, allowing rapid identification of temporal trends, comparison of different types of measurements, and access to detailed information when necessary. As shown in Figure 18, the observation listing interface enables efficient clinical data retrieval.

Patient: JUAN BARBOSA			
Document: 1011320189		Total Observations: 4	
2026-01-19 4 measurements			
16:57 - 17:00	Frecuencia Cardíaca	73.77 beats/min	[N] Normal
16:57 - 17:00	Actividad Física	0.07 MET	[L] Low
16:57 - 17:00	Frecuencia Respiratoria	11.51 breaths/min	[N] Normal
16:57 - 17:00	Postura	23.17 position_code	[A] Abnormal

Figure 18. Observation listing interface.

4.2.4. Observation Detail Visualization Interface

When selecting a specific observation from the list, the detail interface displays all the clinical and technical information contained in the Observation resource generated by PIRE. The detail table includes the LOINC code of the observation with its human-readable description, the status of the observation, the exact start time of the measurement series, the end time of the series, the identifier of the Bioharness-3 device that originated the measurement, the categorical clinical interpretation automatically calculated when processing the observation, the raw measurement data as a statistical summary of the series, the additional information previously entered, and complementary clinical notes recorded during upload. This consolidated view allows clinical professionals, auditors, and administrative staff to fully understand the context, magnitude, and interpretation of each observation without the need to directly access the FHIR server or manipulate technical data structures. As illustrated in Figure 19, the observation detail visualization interface provides comprehensive clinical context.

LOINC CODE	8867-4		STATUS	final
START	2025-09-22T16:57:27.499Z		END	2025-09-22T17:00:40.139Z
DEVICE	Zephyr Bioharness-3		INTERPRETATION	Normal
Measurement Data				
Interval	Minimum Value	Maximum Value		
	0.0 beats/min	93.0 beats/min		
Additional Information				
Method	Monitoreo ambulatorio		Measurement Site	Tórax
Notes				

Figure 19. Observation detail interface.

4.2.5. Observation Deletion Interface

The observation deletion interface allows users with deletion permissions to remove specific measurements from the FHIR repository. To access this functionality, the user enters the identification number of the patient whose observations they wish to manage, and the system displays a list of all observations associated with that patient, similar to the search flow described in section 4.3.3. The user selects the observation to be deleted, confirms the action through a confirmation dialog that specifies the type and date of the observation to be removed, and the application executes an HTTP DELETE request to the HAPI FHIR server to remove the Observation resource from the clinical repository. Once the deletion is completed, the observation disappears from

the list and is no longer available for subsequent queries. As shown in Figure 20, the observation deletion interface enforces authorization controls.

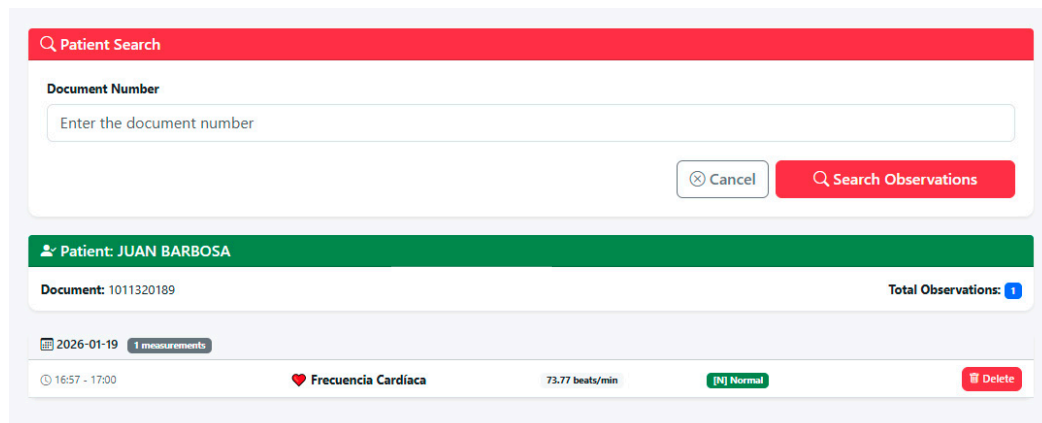


Figure 20. Observation deletion interface.

4.3. User Management

User management in PIRE was implemented as an integrated process that couples the creation of access credentials with the simultaneous generation of Practitioner resources on the HAPI FHIR server, thereby ensuring that each system user has a corresponding clinical representation in the interoperability repository. This dual approach makes it possible to maintain a clear distinction between access identity (username and password) and clinical identity (Practitioner resource), while also guaranteeing an explicit link between both through the storage of the Practitioner identifier in the user record. User management is reserved exclusively for the Super Admin role, which has specialized interfaces to create, view, manage credentials, and delete users, as well as to access detailed information about each professional registered on the platform.

4.3.1. User Creation Interface

The user creation interface in PIRE presents a comprehensive form that guides the Super Admin in capturing all the demographic, contact, and professional information required to establish a new user in the system. The form requests the user's full names, type of identity document, identity document number, and demographic data such as gender and date of birth, information that will later be incorporated into the Practitioner resource generated on the FHIR server. The form continues with the professional's contact data, including phone number, email address, residential address, city, and state or department, allowing the platform to contact the user or geographically locate their clinical practice. Finally, the form requests the specialty or position held by the professional and presents a dropdown list with the three available roles (Doctor, Nurse, Auditor), allowing the Super Admin to explicitly assign the access role that the new user will have on the platform. Once all fields have been completed and validated, the system proceeds to simultaneously create the user record in the application layer and the Practitioner resource on the HAPI FHIR server, automatically generating secure credentials that are presented to the Super Admin to be transmitted to the professional. As illustrated in Figure 21, the user creation interface integrates identity and clinical data capture..

Figure 21. User creation interface.

4.3.2. Practitioner Resource Generation

When a new user is successfully created through the administration interface, PIRE automatically invokes the Practitioner resource generation service, which builds a clinical representation of the professional in accordance with the FHIR R4 specification. The resulting Practitioner resource encapsulates all the demographic, contact, and professional information entered in the creation form, including the unique identifier based on the professional's document, the full name structured into given names and family names, gender, date of birth, phone numbers and email addresses, contact address structured by city and department, and specialty or position of the professional represented through professional classification elements. Once built, the resource is sent to the HAPI FHIR server through an HTTP POST operation directed to the Practitioner endpoint, where it is validated against the FHIR R4 specification and persisted in the server's PostgreSQL database. The unique identifier generated by the FHIR server is captured by the application and stored in the local user record, thus establishing a permanent link between the user's access credential and their clinical representation in the FHIR repository. This integration ensures that all clinical actions performed by the user can be attributed to the corresponding Practitioner on the FHIR server, facilitating the traceability of clinical responsibilities.

4.3.3. "My Profile" Interface

The "My Profile" interface is available to all authenticated users in PIRE, regardless of their role, and provides a consolidated view of the identity, access, and clinical information associated with each user. When accessing this section, the user views their automatically generated username created when the account was created, their password in readable text (only for Super Admin), the role assigned on the platform, the exact date on which the user was created in the system, and basic demographic data such as document type and number, gender, date of birth, professional specialty, and contact information. This view reinforces the transparency of the access model and allows each user to verify the completeness and correctness of their registration information, facilitating the identification of inconsistencies or missing data that may require updating by the Super Admin. As shown in Figure 22, the "My Profile" interface demonstrates user-centered information display.

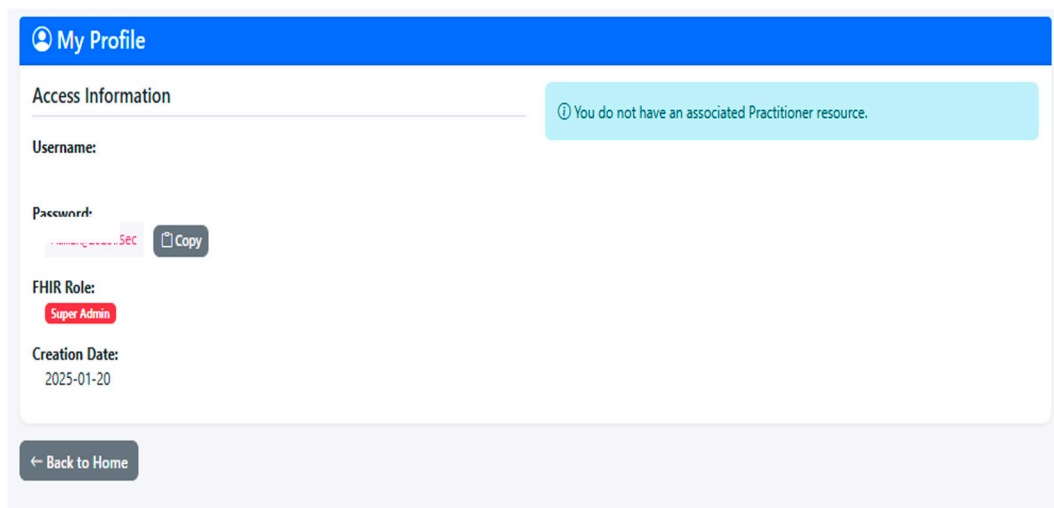


Figure 22. "My Profile" interface.

4.3.4. Password Viewing Interface

The password viewing interface is available exclusively to the Super Admin role and provides access to a consolidated list of all users registered in PIRE with their corresponding passwords in readable text. This interface is designed to facilitate the reissuance of credentials when a user has forgotten or lost their original password, or when it is necessary to provide the credentials again to a professional who requires renewed access to the platform. The Super Admin can view all stored passwords, select the one corresponding to the requesting user, and copy it directly from the interface, without the need to access system files or execute additional technical operations. Once copied, the password can be transmitted to the affected user through a secure communication channel, allowing them to regain access to the platform. The controlled availability of this functionality emphasizes that only the system administrator can recover credentials, limiting the exposure of sensitive information. As illustrated in Figure 23, the password viewing interface restricts credential recovery to authorized administrators.

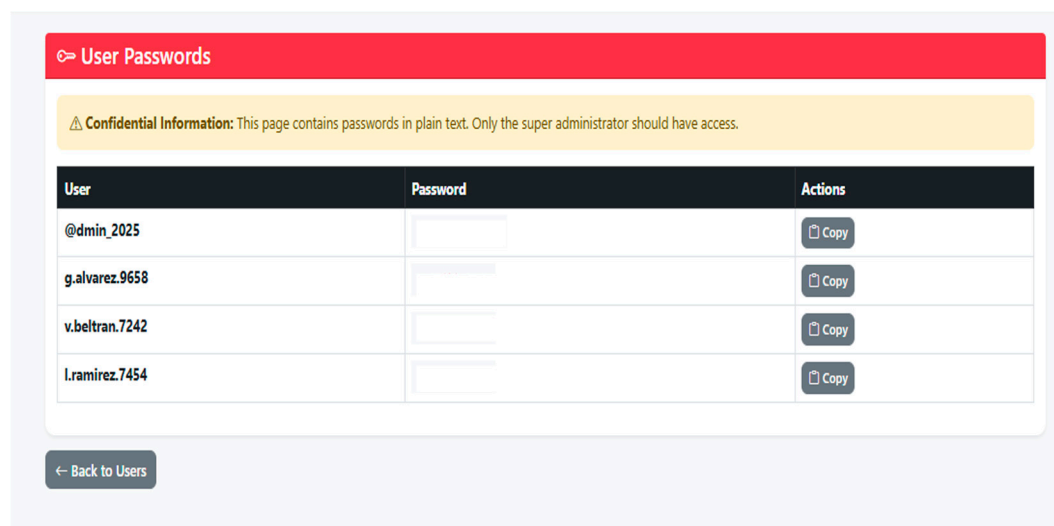
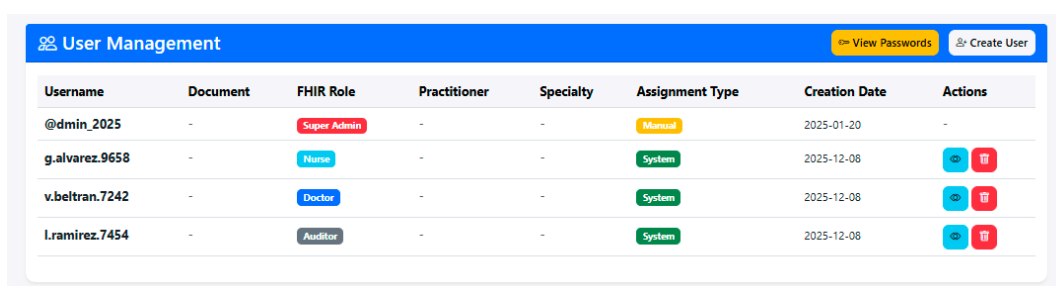


Figure 23. Password viewing interface.

4.3.5. User Listing Interface

The user listing interface presents a consolidated table of all users registered in PIRE with relevant operational information for administration. The table includes columns with the username,

the assigned FHIR role, the type and number of the professional's identity document, the professional specialty or position, the type of role assignment, indicating whether it was assigned through system configuration or manually assigned from the system's internal files (Super Admin), and the exact date of user creation. Each row of the table includes action buttons that allow the administrator to view the detailed information of the user by clicking the view button, which opens an expanded view with complete information. Additionally, each row includes a delete button that allows the Super Admin to remove a user from the system when they are no longer required. When the deletion is executed, the platform implements a cascade process that includes removing the user from the users.csv table, removing their credentials from passwords.csv, deleting the associated Practitioner resource on the HAPI FHIR server, and recording the entire operation in the system audit log, ensuring that no orphan resources remain in the interoperability infrastructure and maintaining full traceability of the user lifecycle. As shown in Figure 24, the user listing interface provides comprehensive user administration capabilities.









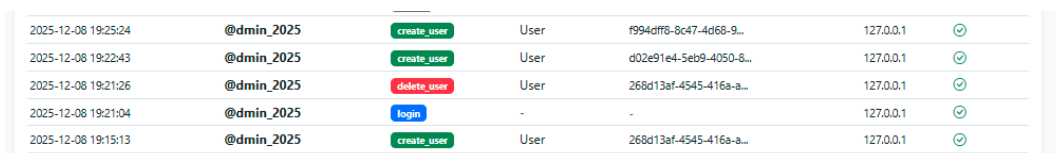
Username	Document	FHIR Role	Practitioner	Specialty	Assignment Type	Creation Date	Actions
@dmin_2025	-	Super Admin	-	-	Manual	2025-01-20	-
g.alvarez.9658	-	Nurse	-	-	System	2025-12-08	 
v.beltran.7242	-	Doctor	-	-	System	2025-12-08	 
l.ramirez.7454	-	Auditor	-	-	System	2025-12-08	 

Figure 24. User listing interface.

4.4. Audit and Logging System

The audit system interface constitutes the traceability layer focused on recording critical events related to user authentication and application login and logout. The interface presents event records in a chronologically ordered table that shows the last one hundred system events. The table allows authorized users to quickly view recent activity, identify specific events, and detect anomalies such as multiple failed login attempts from the same IP address or unusual patterns of user creation and deletion. This view consolidates in a single access point all access and user management traceability information, thereby turning the audit module into an active component of the governance and security model of the PIRE platform. As illustrated in Figure 25, the audit and logging system interface demonstrates comprehensive transaction traceability.



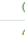


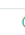
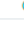
2025-12-08 19:25:24	@dmin_2025	create_user	User	f94dff6-8c47-4d68-9...	127.0.0.1	
2025-12-08 19:22:43	@dmin_2025	create_user	User	d02e91e4-5eb9-4050-8...	127.0.0.1	
2025-12-08 19:21:26	@dmin_2025	delete_user	User	268d13af-4545-416a-a...	127.0.0.1	
2025-12-08 19:21:04	@dmin_2025	login	-	-	127.0.0.1	
2025-12-08 19:15:13	@dmin_2025	create_user	User	268d13af-4545-416a-a...	127.0.0.1	

Figure 25. Audit and logging system interface.

5. Discussion

5.1. Convergence with Historical Evolution and the Colombian Regulatory Framework

The development of the PIRE platform responds to more than six decades of evolution in electronic health records (EHRs), ranging from problem-oriented models in the 1960s to contemporary digital health ecosystems that integrate artificial intelligence and real-time monitoring. By focusing on semantic interoperability, PIRE addresses the most critical challenge identified in the literature: ensuring that clinical information, such as vital signs, preserves its exact meaning when

exchanged, thereby going beyond the technical and syntactic maturity already achieved by other systems [1].

In the national context, PIRE is firmly grounded in the Colombian legal framework, complying with Law 2015 of 2020 by adopting a mandatory model for the exchange of clinical data [34]. Furthermore, the platform integrates the data elements defined by Resolution 866 of 2021 and anticipates the adoption of the Digital Care Summary (Resumen Digital de Atención, RDA) established in Resolution 1888 of 2025, thus enabling local institutions to align with the national interoperability platform and the Vulcano server [35–37]. This strategic alignment is consistent with international experiences, such as those observed in the European Health Data Space (EHDS), where the success of an EHR depends on national architectures that mandate HL7 FHIR-based APIs [8]. In addition, local implementations such as SURAs, which use cloud architectures to harmonize data from more than 10 million patients, empirically validate the feasibility of PIRE by adopting profiles aligned with the FHIR Core CO Implementation Guide [38].

5.2. Technical Architecture and Biosignal Management Using Wearables

The PIRE architecture, based on a HAPI FHIR server deployed on AWS and the use of HL7 FHIR R4, aligns with the view that FHIR is the essential standard for harmonizing biomedical data and enabling clinical decision support tools [10]. In particular, PIRE's ability to process patient-generated health data (PGHD) using the Bioharness-3 device is supported by studies demonstrating the feasibility of integrating continuous metrics into real clinical workflows through FHIR Observation resources. By encapsulating heart rate, respiratory rate, activity, and posture signals using the SampledData structure combined with LOINC terminologies, PIRE ensures that the semantic meaning of each measurement is universally recognizable across interoperable systems. This approach is analogous to implementations that employ Garmin devices or open-source ESP32 microcontrollers to acquire real-time signals such as ECG and PPG, which are subsequently mapped to FHIR resources [28–33]. Moreover, the use of data validators compliant with FHIR Core Colombia, together with a comprehensive logging system, reinforces the integrity of the resources created in PIRE, thereby mitigating the risk of transcription errors or malformed data that could compromise patient safety.

5.3. Integration of Artificial Intelligence and Large Language Models (LLMs)

For future versions of PIRE, the literature suggests powerful tools to automate data normalization and enhance predictive capabilities. While Hou et al. [17] demonstrate that FHIR-Former can predict mortality using structured data, PIRE faces a complementary challenge: whereas Hou et al. focus on structuring clinical narratives, PIRE receives biosignals that are already structured by the Bioharness-3 device. Consequently, the introduction of large language models (LLMs) would be especially valuable not for normalizing vital signs themselves, but for integrating the contextual clinical narratives that accompany these biosignals, for example, correlating increases in heart rate with patient-reported anxiety annotations. In parallel, the adoption of architectures such as FHIR-Former would equip the platform with multimodal predictive capabilities to estimate risks of mortality or hospital readmission based on the combined analysis of structured biosignal data and clinical context [16–18].

These technologies could be integrated into PIRE to overcome the limitations of approaches based solely on manual scripting, enabling the processing of complex clinical narratives that contain ambiguous terminology. Furthermore, the use of voting-based schemes for anomaly detection in biosignals, such as the one proposed by Lee et al., could significantly improve the sensitivity of PIRE's remote monitoring, thereby reducing false negatives in the collected signals [27].

5.4. Security, Governance, and Distributed Access Control

Although Mukherjee et al. [19] propose attribute-based access control (ABAC) to achieve maximum granularity, PIRE currently implements a robust role-based access control (RBAC) model. Nevertheless, a gradual migration toward ABAC would be beneficial as direct data exchange between patients and researchers becomes more common, especially in scenarios where the adoption of smart contracts is critical to ensure that privacy is preserved without relying on intermediaries. To further strengthen distributed security, PIRE could adopt Self-Sovereign Identity (SSI) and blockchain frameworks, thereby enabling patients to become the true owners of their data and facilitating secure access protocols in emergency scenarios [20].

The implementation of smart contracts within environments such as Hyperledger Fabric has proven to be a realistic solution for the transparent and auditable management of large-scale medical records, enabling selective disclosure of information according to the specific needs of verifiers. In this way, PIRE could evolve from a purely technical repository into a secure ecosystem that respects privacy and adheres to the Habeas Data principles mandated by Colombian regulations [20].

5.5. Advanced Analytics and Public Health

PIRE has substantial potential for secondary use of data in research and public health. Specifically, the adoption of the SQL-on-FHIR standard would allow nested FHIR server information to be projected into tabular views, thereby facilitating large-scale statistical analyses by researchers at Universidad Militar Nueva Granada without the need to process complex JSON structures [38]. This capability would open the door to public health studies, such as linking clinical records with geospatial environmental pollution data through pipelines like FHIR PIT [25]. Additionally, the use of Federated Learning (FL) would enable predictive models to be trained directly on patients' devices, preserving full privacy before any sensitive data are sent to the cloud. Such an approach would not only align with data-minimization principles but also position PIRE as a suitable platform for privacy-preserving analytics in distributed clinical environments [28].

5.6. Usability and the Human Factor in Interoperability

A critical aspect that PIRE must address is usability, in order to avoid cognitive overload in clinical staff. Existing studies suggest that the complexity of FHIR can become a barrier to adoption; therefore, the implementation of visual questionnaire editors would help ensure that doctors and nurses manage data without introducing errors [21]. In this regard, PIRE's capability to offer detailed visualizations and value stratifications contributes to reducing cognitive errors during medical consultations by highlighting out-of-range values and simplifying interpretation. At the same time, PIRE explicitly recognizes that interoperability is a socio-technical challenge rather than a purely technical one. Organizational barriers—such as the need for approval by IT leadership and the lack of staff trained in standards such as FHIR—are common obstacles that PIRE must mitigate through continuous training strategies and clear communication of the clinical value of the platform. By doing so, the system can foster not only technical adoption but also organizational engagement [12].

5.7. Divergences and Implementation Lessons

When PIRE is contrasted with other approaches, a clear divergence emerges with respect to systems that depend heavily on legacy standards such as HL7 CDA. Although national and international initiatives such as ELGA in Austria or ENHIS in Estonia demonstrate wide coverage using CDA, the literature increasingly recommends migration toward FHIR to reduce complexity and facilitate dynamic data exchange [5]. In line with Bossenko et al., PIRE recognizes that transformations based solely on manual coding are error-prone; therefore, future implementation of the FHIR Mapping Language (FML), together with visual tools such as TermX, would be preferable to rigid CSV-based processing pipelines [7].

Finally, PIRE deliberately distances itself from implementations that prioritize only technical connectivity. Following socio-technical perspectives, the platform assumes that useful interoperability must focus on the availability and maintainability of records to support real-time clinical decision-making. By integrating commercial devices and continuous signals within a strict regulatory framework, PIRE not only connects institutions but also enhances patient safety and the efficiency of the Colombian health system.

6. Conclusions

The development of PIRE represents a significant milestone in operationalizing semantic interoperability within the Colombian regulatory context. By demonstrating that a fully functional end-to-end FHIR architecture is viable in controlled academic environments, utilizing open-source technologies and explicitly aligned with the Core Colombia profiles, PIRE closes a critical gap between the legislative mandates of Law 215 of 2020, Resolution 866 of 2021, and Resolution 1888 of 2025, and the actual technical capacity of Colombian health institutions to implement them. This alignment is not merely academic, as it validates that medium-sized institutions, which represent the majority of the private health sector in Colombia, can transition toward regulatory compliance without depending exclusively on costly enterprise solutions.

The integration of patient-generated health data through commercial wearable devices such as the Bioharness-3 within a native FHIR architecture significantly expands the scope of what constitutes “interoperable data.” Specifically, PIRE’s capability to encapsulate continuous metrics under standardized LOINC terminologies, while simultaneously preserving the semantics of measurement, resolves a challenge that has not been adequately addressed in prior FHIR literature. This is particularly relevant for remote care models and chronic patient monitoring that are emerging as central components of the ongoing Colombian health system reform.

However, the current validation of PIRE operates within clearly circumscribed limitations. The pilot was conducted in a controlled academic environment with stable network infrastructure, standardized devices, and trained operators—conditions that do not represent the heterogeneity of real clinical settings. The robustness of PIRE in the face of electromagnetic interference, network bandwidth variability, and multiple generations of wearable devices requires additional empirical validation. Equally important, although PIRE currently implements a robust role-based access control (RBAC) model, the evolution toward more granular control models (ABAC) and the integration of self-sovereign identity frameworks are anticipated to become necessary as data exchange becomes decentralized in direct patient-to-researcher or institution-to-institution scenarios.

The incorporation of emerging technologies—such as large language models for processing contextual clinical narratives that accompany biosignals, voting schemes for real-time anomaly detection, and federated learning for training predictive models—constitutes a technically viable roadmap that nonetheless requires rigorous clinical validation. Importantly, these are not merely incremental improvements but rather fundamental transformations in how clinical systems can reason over multimodal data in real-time while preserving privacy and regulatory compliance.

From the perspective of the Colombian health system, PIRE provides three immediate contributions. First, a reproducible reference model that medium-sized institutions can adapt to operationalize IHCE compliance without requiring prohibitive capital investments. Second, technical evidence that clinical semantics can be preserved across data exchange without introducing clinically problematic latency. Third, a demonstration that Colombian regulatory mandates are technically implementable within standard FHIR frameworks, thereby reducing the risk that adoption would require interim solutions that compromise interoperability.

The limitations identified—including validation in controlled environments only, the absence of comparative clinical utility studies, and partial integration of distributed security frameworks—define a clear research agenda for subsequent phases. The transition of PIRE from an academic prototype to an operational platform in pilot institutions would require rigorous clinical validation of the equivalence between wearable measurements and clinical standards, evaluation of socio-

organizational effects on clinical workflows, and security hardening in accordance with international cybersecurity standards for health. These investigations must be conducted with rigor equivalent to that applied in the technical validation.

In conclusion, PIRE not only closes a technical gap in the implementation of semantic interoperability in Colombia, but also repositions the Colombian academic digital health community as a generator of evidence regarding how middle-income health systems can operationalize regulatory transformation in a technically rigorous and economically viable manner. As Colombian health institutions advance along their respective IHCE compliance timelines, PIRE provides both a technical model and an institutional model for how academic research can catalyze health system transformations at the national scale. Ultimately, this work demonstrates that semantic interoperability is not a privilege of wealthy systems but a technically achievable and economically justified goal for health systems committed to evidence-based, standards-driven transformation.

Author Contributions: Conceptualization, LJRL and NEJS; methodology, LJRL; software, LJRL and JEBP; validation, LJRL and NEJS; formal analysis, LJRL and JEBP; investigation, LJRL and NEJS; resources, LJRL; data curation, LJRL; writing—original draft preparation, LJRL, NEJS and JEBP.; writing—review and editing, LJRL; visualization, JEBP; supervision, LJRL and NEJS; project administration, LJRL; funding acquisition, LJRL and NEJS. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Universidad Militar Nueva Granada, grant number INV-ING-4155 and “The APC was funded by LJRL.

Data Availability Statement: Data availability statement: The authors are willing to share the data, codes, and information needed for other researchers.

Acknowledgments: The authors acknowledge the contribution of the researchers from the TIGUM Research Group of the Universidad Militar Nueva Granada of Bogota, Colombia.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ABAC	Attribute-Based Access Control
API	Application Programming Interface
AWS	Amazon Web Services
CDA	Clinical Document Architecture
CIE-10	Clasificación Internacional de Enfermedades, 10ª Edición (International Classification of Diseases, 10th Edition)
CRUD	Create, Read, Update, Delete
CSV	Comma-Separated Values
EHDS	European Health Data Space
ELGA	Elektronische Gesundheitsakte
ENHIS	European Network of Health Information Systems
ESP32	Microcontroller Unit
ETL	Extract, Transform Load
FHIR	Fast Healthcare Interoperability Resources
FHIR CORE CO	FHIR Core Colombia Implementation Guide
FHIR-Former	FHIR-bases Transformer with Large Language Models
FHIRPath	FHIR Path Navigation Syntax
FML	FHIR Mapping Language
Flask	Flask Web Framework
Flask-Session	Flask Session Management Extension
FL	Federated Learning
GIT	Version Control System
HABEAS DATA	Constitutional Right to Data Protection
HAPI	Healthcare API

HAPI FHIR	HAPI FHIR server
HL7	Health Level 7
HL7 CDA	HL7 Clinical Document Architecture
HL7 FHIR R4	HL7 FHIR Release 4
HTML	Hyper Text Markup Language
HTTP	Hyper Text Transfer Protocol
HTTPS	HTTP Secure
IHCE	Interoperable Electronic Health Record (Colombia)
IP	Internet Protocol
IoMT	Internet of Medical Things
IoT	Internet of Things
JSON	JavaScript Object Notation
JPA	Java Persistence API
LOINC	Logical Observation Identifiers Names and Codes
LLM	Large Language Model
Maven	Apache Maven
ONC	Office of the National Coordination for Health Information Technology
OpenEHR	Open Electronic Health Record
Omnisense	Zephyr Omnisense SDK
PDF	Portable Document Format
PGHD	Patient-Generated Health Data
PIRE	Plataforma Interoperable de Registros Electrónicos
POMR	Problem-Oriented Medical Record
PPG	Photoplethysmography
RBAC	Role-Based Access Control
RDA	Resumen Digital de Atención (Digital Care Summary)
REDCap	Research Electronic Data Capture
REST	Representational State Transfer
RMRS	Regenstrief Medical Record System
SDK	Software Development Kit
SFPBRF	SMART on FHIR Platform-Based Reference Framework
SMART	Substitutable Medical Applications and Reusable Technology
SNOMED CT	Systematized Nomenclature of Medicine Clinical Terms
SSH	Secure Shell
SQL	Structured Query Language
SQL Server	Microsoft SQL Server
SURA	Seguros Unidos de Reaseguro Americano
TermX	Terminology Exchange Tool
Ubuntu	Linux Operating System Distribution
VPC	Virtual Private Cloud
WHO	World Health Organization
XOR	Exclusive OR
Zephyr	Zephyr Performance System

References

1. Y. Shen, J. Yu, J. Zhou, and G. Hu, "Twenty-Five Years of Evolution and Hurdles in Electronic Health Records and Interoperability in Medical Research: Comprehensive Review," *J. Med. Internet Res.*, vol. 27, 2025, doi: 10.2196/59024.
2. H. Finnegan and N. Mountford, "25 Years of Electronic Health Record Implementation Processes: Scoping Review," *J. Med. Internet Res.*, vol. 27, 2025, doi: 10.2196/60077.
3. M. H. Gabriel, C. Richwine, C. Strawley, W. Barker, and J. Everson, "Interoperable Exchange of Patient Health Information Among U.S. Hospitals: 2023." [Online]. Available: <https://www.healthit.gov/data/quickstats/electronic-health-information-exchange-hospitals>
4. R. El-Yafouri and L. Klieb, "A scoping review of electronic health records interoperability levels, expectations, approaches, and problems," Oct. 01, 2025, *SAGE Publications Ltd.* doi: 10.1177/14604582251385986.

5. G. Duftschmid, F. Katsch, G. Ciortuz, D. Kalra, and C. Rinner, "Reusing data from HL7 CDA-based shared EHR systems for clinical trial conduct: a method for analyzing feasibility," *BMC Med. Inform. Decis. Mak.*, vol. 25, no. 1, Dec. 2025, doi: 10.1186/s12911-025-02980-2.
6. S. Palojoki, L. Lehtonen, and R. Vuokko, "Semantic Interoperability of Electronic Health Records: Systematic Review of Alternative Approaches for Enhancing Patient Information Availability," 2024, *JMIR Publications Inc.* doi: 10.2196/53535.
7. I. Bossenko, R. Randmaa, G. Piho, and P. Ross, "Interoperability of health data using FHIR Mapping Language: transforming HL7 CDA to FHIR with reusable visual components," *Front. Digit. Health*, vol. 6, 2024, doi: 10.3389/fdgth.2024.1480600.
8. 1, Keyrellous Adib b, 1, Nagui Salama b, 1, Stefania Davia b, Antonio Martínez Millana a, Vicente Traver a, Karapet Davtyan b, * Roberto Tornero Costa a, "Electronic health records and data exchange in the WHO European region: A subregional analysis of achievements, challenges, and prospects," *International journal of medical informatics*, 194, 105687., 2025.
9. Ward Weistra, "The State of FHIR in 2025: Growing adoption and evolving maturity," <https://fire.ly/blog/the-state-of-fhir-in-2025/>.
10. A. Hornback et al., "FHIR in Focus: Enabling Biomedical Data Harmonization for Intelligent Healthcare Systems.," *IEEE Rev. Biomed. Eng.*, vol. PP, Dec. 2025, doi: 10.1109/RBME.2025.3632213.
11. A. Ahmad, F. Azam, and M. W. Anwar, "Implementation of SMART on FHIR in developing countries through SFPBREF," in *ACM International Conference Proceeding Series*, Association for Computing Machinery, Nov. 2018, pp. 137–144. doi: 10.1145/3301879.3301881.
12. A. C. Cheng et al., "Opportunities, barriers, and remedies for implementing REDCap integration with electronic health records via Fast Healthcare Interoperability Resources (FHIR)," *JAMIA Open*, vol. 8, no. 5, Oct. 2025, doi: 10.1093/jamiaopen/ooaf111.
13. J. Kruse, J. Wiedekopf, A. K. Kock-Schoppenhauer, A. Essenwanger, J. Ingenerf, and H. Ulrich, "A Generic Transformation Approach for Complex Laboratory Data Using the Fast Healthcare Interoperability Resources Mapping Language: Method Development and Implementation," *JMIR Med. Inform.*, vol. 12, 2024, doi: 10.2196/57569.
14. M. V. Andersen, I. H. Kristensen, M. M. Larsen, C. H. Pedersen, K. R. Gøeg, and L. B. Pape-Haugaard, "Feasibility of Representing a Danish Microbiology Model Using FHIR," *Stud. Health Technol. Inform.*, vol. 235, pp. 13–17, 2017, doi: 10.3233/978-1-61499-753-5-13.
15. M. Khvastova, M. Witt, A. Essenwanger, J. Sass, S. Thun, and D. Krefting, "Towards Interoperability in Clinical Research - Enabling FHIR on the Open-Source Research Platform XNAT," *J. Med. Syst.*, vol. 44, no. 8, Aug. 2020, doi: 10.1007/s10916-020-01600-y.
16. M. Engelke, G. Baldini, J. Kleesiek, F. Nensa, and A. Dada, "FHIR-Former: enhancing clinical predictions through Fast Healthcare Interoperability Resources and large language models," *Journal of the American Medical Informatics Association*, vol. 32, no. 12, pp. 1793–1801, Dec. 2025, doi: 10.1093/jamia/ocaf165.
17. Z. Hou, M. Jiang, H. Liu, and Y. Zhuang, "LLM-Integrated Normalization and Knowledge for FHIR (LINK-FHIR)," in *Studies in Health Technology and Informatics*, IOS Press BV, Aug. 2025, pp. 17–21. doi: 10.3233/SHTI250793.
18. J. Grimes et al., "SQL on FHIR - Tabular views of FHIR data using FHIRPath," *NPJ Digit. Med.*, vol. 8, no. 1, Dec. 2025, doi: 10.1038/s41746-025-01708-w.
19. S. Mukherjee, I. Ray, I. Ray, H. Shirazi, T. Ong, and M. G. Kahn, "Attribute based access control for healthcare resources," in *ABAC 2017 - Proceedings of the 2nd ACM Workshop on Attribute-Based Access Control, co-located with CODASPY 2017*, Association for Computing Machinery, Inc, Mar. 2017, pp. 29–40. doi: 10.1145/3041048.3041055.
20. A. L. Martínez, M. Naghmouchi, M. Laurent, J. G. Alfaro, M. G. Pérez, and A. R. Martínez, "Breaking barriers in healthcare: A secure identity framework for seamless access," *Comput. Stand. Interfaces*, vol. 95, Jan. 2026, doi: 10.1016/j.csi.2025.104020.

21. C. Vogel, R. Pryss, P. Heuschmann, V. Rücker, and M. Winter, "Assessing a visual editor for healthcare questionnaires based on the fast healthcare interoperability resources (FHIR) standard: protocol for a cross-sectional, mixed-methods usability evaluation using eyetracking and retrospective think-aloud," *BMJ Open*, vol. 15, no. 9, Sep. 2025, doi: 10.1136/bmjopen-2025-099744.
22. N. Haffer et al., "Genomics on FHIR – a feasibility study to support a National Strategy for Genomic Medicine," *NPJ Genom. Med.*, vol. 10, no. 1, Dec. 2025, doi: 10.1038/s41525-025-00516-1.
23. C. S. Schmidt et al., "A real-time dashboard for optimizing blood product utilization through a FHIR-based data integration pipeline," *Transfusion and Apheresis Science*, vol. 65, no. 1, Feb. 2026, doi: 10.1016/j.transci.2025.104301.
24. C. Anderson, M. Algorri, and M. J. Abernathy, "Real-time algorithmic exchange and processing of pharmaceutical quality data and information," Oct. 15, 2023, *Elsevier B.V.* doi: 10.1016/j.ijpharm.2023.123342.
25. K. Fecho, J. J. Garcia, H. Yi, G. Roupe, and A. Krishnamurthy, "FHIR PIT: a geospatial and spatiotemporal data integration pipeline to support subject-level clinical research," *BMC Med. Inform. Decis. Mak.*, vol. 25, no. 1, Dec. 2025, doi: 10.1186/s12911-024-02815-6.
26. A. Cruoglio et al., "An Extract, Transform, Load foundation for Biobank Data Interoperability," *Journal of Data and Information Quality*, vol. 17, no. 4, Dec. 2025, doi: 10.1145/3769117.
27. J. Lee, J. Jin, and Y. Kim, "Voting-based real-time monitoring scheme for healthcare data on medical record clouds," in *2024 Research in Adaptive and Convergent Systems - Proceedings of the 2024 International Conference on Research in Adaptive and Convergent Systems, RACS 2024*, Association for Computing Machinery, Inc, Oct. 2025, pp. 157–159. doi: 10.1145/3649601.3698739.
28. A. Akhmetov, Z. Latif, B. Tyler, and A. Yazici, "Enhancing healthcare data privacy and interoperability with federated learning," *PeerJ Comput. Sci.*, vol. 11, pp. 1–32, 2025, doi: 10.7717/peerj-cs.2870.
29. F. A. Seixas-Lopes, C. Lopes, M. Marques, C. Agostinho, and R. Jardim-Goncalves, "Musculoskeletal Disorder (MSD) Health Data Collection, Personalized Management and Exchange Using Fast Healthcare Interoperability Resources (FHIR)," *Sensors*, vol. 24, no. 16, Aug. 2024, doi: 10.3390/s24165175.
30. P. Randine and E. Årsand, "FHIR-Based Information System for Type 1 Diabetes Consultations Based on Patient-Generated Health Data," in *Studies in Health Technology and Informatics*, IOS Press BV, Aug. 2025, pp. 189–193. doi: 10.3233/SHTI250827.
31. P. Urbauer, V. David, M. Frohner, and S. Sauermann, "Wearable activity trackers supporting elderly living independently: A standards based approach for data integration to health information systems," in *ACM International Conference Proceeding Series*, Association for Computing Machinery, Jun. 2018, pp. 302–309. doi: 10.1145/3218585.3218679.
32. S. Abedian, E. Yesakov, S. Ostrovskiy, and R. Hussein, "Streamlining wearable data integration for EHDS: a case study on advancing healthcare interoperability using Garmin devices and FHIR," *Front. Digit. Health*, vol. 7, 2025, doi: 10.3389/fdgth.2025.1636775.
33. F. C. Adochiei, F. A. Țoi, I. R. Adochiei, F. C. Argatu, G. Serîțan, and G. G. Petroiu, "HL7 FHIR-Based Open-Source Framework for Real-Time Biomedical Signal Acquisition and IoMT Interoperability," *Applied Sciences (Switzerland)*, vol. 15, no. 23, Dec. 2025, doi: 10.3390/app152312803.
34. Ministerio de Salud y Protección Social de Colombia, "Ley 2015 de 2020," 2020.
35. Ministerio de Salud y Protección Social de Colombia, "Resolución 866 de 2021," 2021.
36. MINISTERIO DE SALUD Y PROTECCIÓN SOCIAL, "MINISTERIO DE SALUD Y PROTECCIÓN SOCIAL," 2025.
37. Ministerio de Salud y Protección Social de Colombia, "Resumen Digital de Atención en Salud (RDA)," <https://vulcano.ihcecol.gov.co/>.
38. M. Nelson and J. Hausam, "Global Use of HL7 Standards Driving Better Interoperability PLUS: Year-end Recaps from Vulcan and FAST First Gravitare-Health Hackathon Generates Solutions Recap of HL7 Europe WGM 2024 and FHIR Marathon Interoperability Project to Improve Healthcare Services in One of the Largest Insurance Companies in Colombia CMS Encourages Use of HL7 Da Vinci Project's FHIR API Implementation Guide to Comply with Rule INSIDE: HL7 NAMES FOUR AS 2023 VOLUNTEERS OF THE YEAR."

39. MEDIFOLIOS, "Asistene de registros médicos," <https://www.medifolios.net/modulos/integracion-ecopetrol.php>.
40. MEDIFOLIOS, "HCE: Normativas y Beneficios Colombia," <https://medifolios.net/articulos/historia-clinica-normativas-beneficios.php>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.