

Chapter 7 / Capítulo 7

Primum non nocere: Software engineering
for medicine
ISBN: 978-9915-9876-8-2
DOI: 10.62486/978-9915-9876-8-2.Ch07
Pages: 111-132



© 2026 The authors. This is an open access
book chapter distributed under the terms of
the Creative Commons Attribution (CC BY)
4.0 License.

Chapter 7. Interoperability Standards

Annier Jesús Fajardo Quesada  

Eduardo Antonio Hernández González 

René Herrero Pacheco 

Rolando Javier Álvarez Pérez 

1 Granma University of Medical Sciences. Granma, Cuba.

2 University of Medical Sciences of Pinar del Río. Pinar del Río, Cuba.

“Interoperability is neither a technical luxury nor a distant goal: it is the minimum requirement for 21st-century medicine to be practiced reliably. When systems cannot communicate with one another, the one who pays the price is not the IT specialist who failed to make them talk, but the patient whose medical history is fragmented, the physician who makes decisions with incomplete information, and the healthcare system that multiplies unnecessary tests.”

Imagine a patient visiting a cardiology clinic. The cardiologist needs to know the results of the latest lab tests, discharge summaries from previous hospitalizations, the current medications prescribed by the patient's primary care physician, and imaging tests performed at another facility. In an ideal world, all this information would be available on the computer screen and integrated in a coherent, up-to-date, and reliable manner. In the real world, what usually happens is very different: the cardiologist has to ask the patient to bring the reports, call the health center, or ask the patient to repeat the tests because the previous results are not accessible. Behind this everyday scenario lies a problem that is not medical but technical, and it has a name: lack of interoperability.

Interoperability is the ability of different information systems to communicate, exchange data, and use exchanged information. In the healthcare sector, this ability is not an option but a clinical necessity. Patients are not treated in a single location by a single doctor using a single system. They move between primary and specialty care, between emergency rooms and hospitalization, between public and private facilities, between autonomous communities, and sometimes between countries. For their information to follow them, for every professional treating them to have access to the relevant data, the systems must be able to understand one another. It is not a matter of administrative efficiency but of clinical safety.

In this chapter, we explore the main standards that enable interoperability in healthcare. We start with the FHIR, the most modern and promising standard, which is revolutionizing the way healthcare applications exchange information. We will also look at HL7 v2, the legacy standard that still dominates many hospital systems, and DICOM, the standard that enables the exchange of medical images. Finally, we explore SMART on the FHIR, an extension that enables the development of applications that integrate directly into electronic health records systems. This is not an exhaustive technical guide but rather an overview designed to help clinicians and developers understand the possibilities and limitations of each standard and make informed decisions about when and how to use them.

7.1 FHIR: PHILOSOPHY, RESOURCES, AND USE CASES

At the end of the last decade, the world of healthcare interoperability found itself in a frustrating paradox. Technical standards existed—some with decades of history and maturation—but they were complex to implement, rigid in their architecture, and difficult for modern developers to adopt, given their familiarity with elegant REST APIs, agile data exchange in JSON, and rapid iterative development cycles. Integration projects between healthcare systems inevitably dragged on for months or years, required highly specialized and expensive consulting experts, and often resulted in highly customized solutions that were neither reusable nor maintainable. Against this backdrop of widespread frustration, a group of visionaries within HL7 International—the organization historically responsible for healthcare interoperability standards—decided that it was time to start over with a radically different approach. Not simply to create yet another technical standard among many but to rethink from scratch what the exchange of health information should look like in the modern era of the internet and mobile computing. The result of that effort was FHIR, and it literally changed the game for healthcare interoperability.⁽¹⁾

FHIR, which stands for fast healthcare interoperability resources, is an open standard for the exchange of healthcare information that combines the best of previous semantic standards with modern web technologies that developers already know and master. Its name says it all: it is fast because it uses stateless REST APIs and lightweight exchange formats such as JSON; it is specifically for healthcare because it is designed from the ground up for the clinical domain with all its complexities; and it is fundamentally based on resources, which are the fundamental building blocks upon which the entire interoperability architecture is built. Since its official launch, the FHIR has been widely adopted by leading commercial electronic health record vendors (Epic, Cerner, and Meditech), by government health agencies around the world, and by thousands of independent developers who have found in it a finally simple and accessible way to make

healthcare systems talk to each other.⁽¹⁾

7.1.1 THE THREE FUNDAMENTAL PRINCIPLES OF THE FHIR

FHIR's design philosophy is based on several fundamental principles that clearly distinguish it from previous interoperability standards. The first principle is intentional simplicity: The FHIR is explicitly designed to be easy to understand and implement, even for software developers without specific prior experience in healthcare interoperability. It uses technological concepts familiar to any modern web developer: RESTful architecture with standard HTTP verbs (GET, POST, PUT, and DELETE), resource exchange in JSON or XML depending on preference, authentication via OAuth 2.0, and interactive online documentation. This technological familiarity drastically reduces the historically steep learning curve and makes integration between systems much faster to implement and less costly to maintain.

The second principle is architectural modularity. The FHIR is not a rigid, indivisible monolith but rather a set of more than 150 independent and well-defined resources that can be used separately or combined according to the specific needs of the use case. A resource can conceptually represent a patient, an individual clinical observation, a medication prescription, a diagnostic condition, a performed procedure, or any of the many types defined in the standard. Each resource has a well-defined data structure, with mandatory fields for identification and optional fields for semantic richness, and can be formally extended via profiles to adapt to specific local needs without breaking core compatibility with other FHIR systems.

The third principle is evolution based on the best practices of the modern web. FHIR explicitly adopts the same software engineering practices that have made the global web successful and scalable: semantic versioning of specifications (major.minor.patch), guaranteed backward compatibility within major versions, interactive documentation with executable examples, and an active global community that openly shares reference implementations, national profiles, and lessons learned. This means that, unlike previous standards that require costly, disruptive updates every few years, the FHIR can evolve gradually and continuously, adding new resources and functionalities without intentionally breaking what is already working in production.

Resources are the architectural heart of the FHIR and its fundamental unit of exchange. A resource is a structured, typed representation of a clinical, administrative, or logistical concept in the healthcare domain that is explicitly designed to be exchanged between heterogeneous systems. Each resource has a globally unique identifier (ID) and a set of typed data elements (such as name, date, numeric value, terminology code) and may contain formal references to other related resources. For example, an "observation" resource representing a systolic and diastolic blood pressure measurement may reference the "patient" resource representing the patient on whom the measurement was taken, the "encounter" resource representing the clinical encounter during which the measurement was taken, and the "practitioner" resource representing the healthcare professional who performed or validated the measurement. This network of references between resources allows the complete clinical context to be computationally reconstructed from individual exchanged resources.

7.1.2 THE RICHNESS OF RESOURCES AND THE CONCEPT OF A PROFILE

The list of standard FHIR resources covers virtually the entire semantic domain of modern healthcare practice. There are fundamental resources for demographic and organizational data (Patient, Practitioner, PractitionerRole, Organization, Location), for clinical process management (Encounter, EpisodeOfCare, Appointment, Schedule, Slot), for recording diagnoses and conditions (Condition, AllergyIntolerance, FamilyMemberHistory, ClinicalImpression), for documenting procedures and interventions (Procedure, MedicationAdministration, Immunization, CarePlan), for observations and results (Observation, DiagnosticReport, ImagingStudy, Media, Specimen), for medication management (Medication, MedicationRequest, MedicationDispense, MedicationStatement), and many others for administrative, financial, and follow-up functions. This broad and deep semantic coverage enables the FHIR to be used in virtually any conceivable interoperability scenario, ranging from simple document exchange between organizations to real-time integration of IoT medical devices, including structured access to data for clinical research or artificial intelligence.⁽¹⁾

One of the most important and powerful innovations of the FHIR is the formal concept of

a “profile.” Since the base standard must be sufficiently flexible and generic to adapt to different national contexts and specialties, the base resources have many optional elements and possible extensions. A profile is a formal restriction of those base resources for a specific use, defining which elements are mandatory in that context, which terminology vocabularies must be used for coding, which local extensions are necessary, and which value restrictions apply. For example, a Spanish national profile may define how patient identification numbers (DNI, NIE, health card) must be specifically represented in the patient resource. A profile for a specific clinical trial may define which outcome observations must be collected and in what standardized format. Profiles allow the FHIR to be simultaneously a globally interoperable standard and a tool locally adapted to specific needs without fragmentation.

7.1.3 TRANSFORMATIVE USE CASES AND THE FUTURE OF FHIR

The practical use cases for FHIRs are as varied and numerous as the global healthcare practice itself is. One of the most common and visibly transformative is access to patient data via consumer mobile apps. A patient can use an app on their smartphone to securely view their recent lab results, current active medications, upcoming scheduled appointments, or even download their complete medical record, provided that their hospital or insurance provider's system exposes that data through an authenticated and authorized FHIR API. This same API can be used by a primary care physician to access patient data from their own certified application or by an authorized academic researcher to extract aggregated, anonymized data for a population-based epidemiological study.

Another high-impact use case is the seamless integration of medical devices and health wearables. A continuous glucose monitor, a connected digital blood pressure monitor, a smart pacemaker, or an IoT bathroom scale can automatically send their readings to the patient's electronic health record using standard FHIR resources (typically Observation). This eliminates the need for patients or healthcare professionals to manually enter data, reducing transcription errors, improving the quality of information with data of greater temporal granularity, and facilitating the remote monitoring of chronic patients. The accelerating trend toward connected medicine and hospitals without walls, with devices that are increasingly smart, ubiquitous, and affordable, makes this use case increasingly strategic for future care models.

Clinical research is another area where the FHIR has a significant effect on data efficiency and quality. Traditional clinical trials require the collection of large amounts of structured data from participants, and this was traditionally done using paper-based data collection notebooks or separate electronic data capture (EDC) systems disconnected from actual electronic health records. With FHIR, it is technically possible to automatically extract relevant data from the participant's electronic health record, drastically reducing the burden of double data entry, minimizing transcription errors, and improving data quality and accuracy since it comes directly from the routine clinical record. Initiatives such as the U.S. ONC's Sync for Science project are exploring precisely this transformative possibility, allowing patients to authorize the sharing of their structured FHIR data with researchers in a simple, secure, and ethically supervised manner.

Continuity of care across levels of care is perhaps the most classic use case, yet one that remains unsatisfactorily resolved in many systems. When a patient is referred from primary to specialty care or when they are discharged from the hospital to their home, it is crucial that relevant clinical information travels with them in a structured and actionable format. FHIRs allow this information to be packaged into a “DocumentReference” or a “Bundle” containing all relevant resources: coded diagnoses, active medications with dosages, documented allergies, tests performed with results, and a discharge care plan. This “package can be transmitted electronically to the receiving system, which can incorporate it into its own repository while maintaining the structure, clinical meaning, and traceability of the data.

For software developers, working with the FHIR is now much more accessible and productive than working with previous standards. There are mature and well-documented client libraries for virtually all modern programming languages (Java, C#, Python, JavaScript/TypeScript, and Swift), open-source FHIR servers (HAPI FHIR, Smile CDR, and Microsoft FHIR Server) that can be deployed locally for development and testing, and an active global community on forums and GitHub that shares experiences, resolves technical questions, and contributes improvements. The official HL7 FHIR documentation is clear, wellorganized by resources and operations, and includes interactive, executable examples that allow you to test the APIs directly from your browser without the need to install any local software. This technological accessibility has been key to the rapid adoption of

the FHIR, especially among innovative startups and small developers who previously could not afford the cost of entry into the closed world of healthcare interoperability.

For physicians and clinical professionals involved in FHIR system development or implementation projects, understanding the basic principles of this technology helps them grasp the new possibilities it opens up for improving clinical practice. They can envision and propose clinical applications that previously seemed impossible or prohibitively expensive: an intelligent prescription assistant that automatically checks for drug interactions with the patient's current medications in real time via FHIR; a predictive alert system that notifies staff when a high-risk patient has missed scheduled follow-up appointments; and an integrated visualization tool that aggregates data from multiple sources (hospital, laboratory, pharmacy, devices) into a single longitudinal monitoring screen. All of these applications are technically possible today with FHIRs, and many of them are already in use at leading hospitals around the world.

Information security is a fundamental and nonoptional aspect of FHIRs. Since the APIs expose highly sensitive clinical data protected by strict regulations (HIPAA, GDPR), the standard incorporates robust authentication and authorization mechanisms based on industry standards such as OAuth 2.0 for authorization and OpenID Connect for identity authentication. Patients can granularly control which third-party applications access which categories of their data and for how long, revoking access whenever they wish. FHIR systems can comprehensively audit every data access to detect potential misuse or intrusions. This security layer, combined with TLS encryption in transit and at rest, makes FHIR technically suitable even for the most demanding and heavily regulated healthcare environments.

The future of the FHIR as the dominant standard is promising and increasingly well-established. The leading electronic health record systems in the global market (Epic, Cerner, Meditech, Allscripts) already offer certified FHIR APIs in their products, and many countries have formally adopted it as a mandatory national standard for interoperability. The U.S. Office of the National Coordinator for Health Information Technology (ONC) has mandated the use of the FHIR for the certification of new health systems, which has dramatically accelerated its adoption in the North American market. In Europe, the European Health Data Space (EHDS) is actively promoting the use of the FHIR for the cross-border exchange of patient summaries and research health data. All signs point to the FHIR becoming the de facto universal common language of digital health in the coming years—the technical Esperanto that will allow systems worldwide to communicate effectively with one another.

It is important to remember, realistically, that the FHIR is not a magic solution that automatically resolves all historical interoperability problems in healthcare. Careful and expert work remains absolutely necessary in architectural design, and appropriate semantic profiles for each context, quality data governance, and organizational change management for adoption must be defined. However, the FHIR has drastically lowered the technical barrier to entry, ultimately making interoperability accessible to a much broader and more diverse community of developers, entrepreneurs, and clinical innovators. In addition, in a world where technology is advancing ever faster and the need for coordinated care is becoming increasingly urgent, it is likely its most valuable and lasting contribution to improving global health.

7.2 HL7 V2: THE LEGACY THAT WON'T GO AWAY

In any hospital we visit today, we'll find a reality that surprises developers new to the healthcare sector. Systems that have been in operation for twenty or thirty years, built on technologies that are archaic by today's standards, remain the beating heart of daily operations. Their communication interfaces are based on a standard born in the 1990s, when the internet was not yet the global infrastructure we know today. Every time someone proposes modernizing it, the same inevitable dilemma arises: replacing everything would be economically prohibitive and clinically risky, but keeping it means living with limitations that hinder innovation. This hospital is not an exception but the absolute norm in healthcare across half the world.

In terms of all the odds, HL7 v2 remains the dominant standard for systems integration in hospital settings. Everything indicates that it will continue to be so for many years to come, not for lack of modern alternatives, but owing to the inertia of thousands of critical implementations that no one can afford to shut down. Understanding why this standard persists and understanding how to coexist with it intelligently has become an essential skill for any developer aspiring to work in real-world healthcare environments. Ignoring its existence is equivalent to attempting to build bridges without knowing the ground on which the foundations rest.

7.2.1 THE ORIGINS OF A PRAGMATIC STANDARD

HL7 v2 was born in the late 1980s and early 1990s, when connectivity was a luxury rather than a commodity. The internet did not yet exist as a global infrastructure, and hospitals needed to exchange information between systems that rarely shared the same manufacturer. The solution proposed by Health Level Seven International was surprisingly simple: plain text messages, machine-readable but also interpretable by humans with some training, based on simple delimiters such as the pipe (|) and the caret (^). These messages could be transmitted over basic protocols such as TCP/IP or even serial connections, adapting to the technical limitations of the time without the need for sophisticated infrastructure.(2)

The design philosophy reflected the needs of an era where bandwidth efficiency was crucial and processing power was limited. What might seem like a cryptic and unnecessarily complex monstrosity to a modern developer, accustomed to REST APIs with JSON and OAuth authentication, represented an elegant and relatively simple solution to implement for the systems of that era. It requires no heavy libraries, no proprietary binary formats, and no complicated external dependencies. A program written in COBOL could generate these messages just as easily as one written in C or in the early fourth-generation languages.

The structure of an HL7 v2 message follows conventions that, once understood, are logical although verbose. It typically begins with an MSH (message header) segment that serves as the message's cover letter, indicating who is sending it, who should receive it, what type of message it contains, and when it was generated. This header is followed by other specialized segments depending on the purpose of the communication: PID segments for patient demographics, PVI for information about the care episode, OBR for diagnostic test orders, and OBX for the results themselves. Each segment consists of fields separated by pipes, and each field can be broken down into subfields using circumflex accents.

The visual result is a long and seemingly cryptic text string, where the information appears densely packed without the whitespace or indentation that make other formats readable. However, this apparent chaos follows strict rules and is perfectly interpretable by any system that knows the grammar. A laboratory results message may contain dozens of OBX segments, each with multiple fields describing everything from the test code to the numerical value obtained, including units of measurement, reference ranges, and technician's notes. The density of information is remarkable, even if human readability is not desired.

What made HL7 v2 a successful standard—and what explains its surprising longevity—was precisely its combination of simplicity and flexibility. Unlike more rigid standards that attempted to anticipate and standardize every possible scenario from the outset, HL7 v2 adopted a pragmatic and modest approach. It provided a robust basic structure but allowed organizations to adapt it to their local needs through extensions, optional fields, and bilateral implementation agreements. This flexibility made it relatively easy to implement even for manufacturers with limited resources.

The ability to adapt to the enormous diversity of systems and processes existing in hospitals around the world was decisive. Each facility had its own unique legacy workflows and local coding systems. A standard that was too rigid would have failed in the face of this variability, but HL7 v2 managed to become the de facto standard for integration because it knew how to yield to local realities without compromising global interoperability. As a result, during the 1990s and 2000s, it became the lingua franca of communication between hospital systems—the dialect that everyone understood, even though each spoke it with their own accent.

7.2.2 CURRENT REALITY: PERSISTENCE AND DOMINANCE

Today, decades after its initial creation, HL7 v2 maintains a dominant position that baffles those who predicted its demise on multiple occasions. Laboratory systems continue to send results to electronic health records via HL7 v2 messages, exactly as they did in the 1990s. Hospital pharmacy systems receive electronic prescriptions in the same format. Admission systems report admissions, discharges, and transfers using the same syntax as always. Approximately ninety percent of U.S. hospitals use HL7 v2 for some of their critical integrations, and the situation in Europe, Latin America, and other regions does not differ substantially.

This persistence is not due to a lack of modern alternatives. The FHIR, the most recent standard from the same organization, HL7 International, offers technically superior capabilities in almost every respect. However, the economic and operational realities of hospitals create formidable inertia. Replacing a laboratory system that works, is validated, is thoroughly familiar

to staff, and is integrated with dozens of other systems via HL7 v2 represents a multimillion-dollar cost and unacceptable operational risk. Even quite rightly, hospital executives wonder why they should risk the continuity of care to obtain technical benefits that often do not translate into perceptible clinical improvements.

For developers new to the healthcare field, working with HL7 v2 can be a deeply bewildering experience. Messages are difficult for a human to read without specific tools; they lack the visual hierarchical structure offered by formats such as XML or JSON, and the available documentation can be overwhelming because of its length and technical density. Debugging tools are limited compared with modern development ecosystems, and the testing cycle often requires complex environments that simulate the hospital setting. However, it is an absolutely necessary skill because most real-world integrations in hospitals involve HL7 v2 at some point in the process.

Ignoring this standard is tantamount to condemning oneself to work in a technological fantasy world where only modern, well-designed systems exist when the reality of healthcare is radically different. A developer who proposes a solution based solely on modern REST APIs, without considering how it will integrate with the legacy systems that handle eighty percent of clinical data, is building on quicksand. Hospitals are not software labs where systems can be rebooted every year; they are living organisms where continuity is a priority and where every change must be clinically justified.

One of the most important features of HL7 v2, and one of its greatest sources of operational complexity, is the multitude of versions coexisting in the ecosystem. There are HL7 v2.1, v2.2, v2.3, v2.3.1, v2.4, v2.5, v2.5.1, v2.6, v2.7, and intermediate versions that add specific functionalities. Each version incorporates new segments, new fields within existing segments, and new functionalities that reflect the evolution of clinical practice. However, the most problematic issue is not the existence of multiple versions but rather that not all systems implement them consistently or completely.

Often, two systems that supposedly speak the same version of HL7 v2 have different interpretations of how certain optional fields should be used or of which values are valid for specific codes. This variability means that every integration between two specific systems is a custom engineering project, requiring detailed analysis of each system's particular specifications and exhaustive testing to ensure that communication is semantically correct. It is not enough for the message to be syntactically valid; it must convey the correct clinical meaning between systems that may have radically different internal data models.

7.2.3 COEXISTENCE STRATEGIES AND COMPLEXITY MANAGEMENT

To manage this inherent complexity, most medium and large hospitals have adopted an architecture based on integration engines, also known as interface engines. These specialized products act as intelligent intermediaries between different systems, receiving messages in one format, translating them into another when necessary, routing them to the correct recipients according to configurable rules, and providing monitoring and debugging tools that compensate for the limitations of the original systems. A typical integration engine can receive an HL7 v2 message from a laboratory, convert its local coding to a hospital standard, enrich it with additional information from other systems, and deliver it to the electronic health record in the exact format it expects.

These engines are, in reality, the invisible glue that holds together the ecosystem of heterogeneous systems in a modern hospital. Without them, each system would have to know the specific details of every other system with which it communicates, creating a network of dependencies that would be impossible to manage. With a centralized integration engine, each system only needs to know how to communicate with the engine, and it is the engine that handles the complexities of translation and routing. The most established products on the market, such as Mirth Connect (now NextGen Connect), Rhapsody, or Intersystems Ensemble, have perfected these capabilities over decades.

The relationship between HL7 v2 and more modern standards such as the FHIR is not, as is sometimes portrayed, one of confrontation or imminent replacement. The FHIR is not going to “kill” HL7 v2, at least not within the foreseeable timeframe of the next decade. What is happening in practice is a structured coexistence where each standard occupies its functional niche. Hospital systems are exposing FHIR APIs for new developments, mobile applications, and patient portals, while their critical internal integrations and connections to legacy systems continue to use HL7 v2 exactly as before.

Integration engines are evolving into bilingual bridges, adding native capabilities to translate between both worlds. In this way, a modern FHIR-based application can communicate seamlessly with a 1990s-era lab system that only understands HL7 v2, without either end needing to know the other exists. This translation architecture is the practical reality of contemporary healthcare interoperability and will likely remain so for many years, given the pace of hospital system upgrades.

For developers building new applications for the healthcare market, a sensible architectural strategy is to use the FHIR to interface with the outside world and with modern systems but to be prepared to integrate with systems that speak only HL7 v2 when necessary. This involves understanding the basics of the legacy standard, knowing how to read and interpret messages manually when automated tools fail, and understanding how integration engines work to collaborate effectively with hospital infrastructure teams. It also means being aware of the inherent security limitations.

HL7 v2 was designed at a time when cybersecurity was not the central concern it is today. The original standard lacks a robust security model comparable to that of FHIRs and does not support modern token-based authentication, and messages often travel in plain text unless additional encryption measures are taken at the transport layer. Integrating legacy systems requires special attention to security and privacy and implementing encrypted tunnels, VPNs, or network segmentation to compensate for the application protocol's shortcomings. Clinical data are sensitive and contain regulated information, and we cannot allow them to travel across hospital networks without adequate protection.

For healthcare professionals involved in development projects, understanding that this legacy world exists helps them maintain realistic expectations regarding the timelines and complexity of integration. When a developer explains that accessing lab data will take weeks of work because they must understand exactly how the current system sends it, they are not making excuses or exaggerating difficulties but rather accurately describing the technical reality. Legacy systems lack well-documented APIs, do not uniformly follow standards, and often require reverse engineering to determine what information the messages they generate actually contain.

One frequently overlooked aspect is that HL7 v2, despite its age, remains the subject of active maintenance and controlled evolution. HL7 International regularly publishes new versions that incorporate requirements from contemporary clinical practice, although the changes are necessarily incremental and maintain backward compatibility so as not to break existing implementations. There are also profiles and implementation guides that help reduce problematic variability between systems. In the United States, the IHE (Integrating the Healthcare Enterprise) initiative defines specific usage profiles for HL7 v2 in specific scenarios such as the exchange of laboratory reports or the notification of hospital discharge.

These profiles facilitate integration by reducing the available options and establishing common conventions on how to use optional fields, which encodings to employ, and how to handle exceptional scenarios. When two systems follow the same IHE profile, the likelihood of seamless interoperability increases significantly. However, the adoption of these profiles is not universal, and many legacy systems were developed before they existed or implemented them only partially. The quality of HL7 v2 implementations is, consequently, extremely variable in the market.

Some systems, especially the most modern ones or those developed by large corporations, follow the specifications of the letter and generate clean, well-structured, and predictable messages. Others, particularly those developed in-house by hospitals or by small local vendors, have highly proprietary implementations, with fields used in undocumented ways, local encodings without standard equivalents, and erratic behaviors that defy all logic. Part of the daily integration work consists precisely of addressing this variability, meticulously documenting the peculiarities of each system and developing adapters that normalize messages to a common internal format.

Monitoring HL7 v2 communications is another critical aspect that we cannot underestimate. In a hospital environment where information must flow continuously to support clinical decisions, messages can be lost because of network errors, corrupted by encoding issues, or arriving out of order because of variable latencies. Systems can go down because of maintenance or failures, resulting in unprocessed messages in overflowing queues. This is why it is essential to have tools that allow for real-time visualization of message flow, detection of statistical anomalies, alerts for interruptions, and reprocessing of messages that were not delivered correctly.

Integration engines typically include basic monitoring capabilities, but there are also specialized tools that provide comprehensive visibility to the entire communication ecosystem, allowing technical teams to identify bottlenecks, analyze traffic patterns, and anticipate problems

before they impact clinical operations. The availability of these monitoring systems is just as important as the integration capability itself because in a hospital, a silent interface can mean that a physician is not receiving critical lab results.

Looking ahead, HL7 v2 will likely continue to exist for decades, simply because the cumulative cost of replacing all systems that use it is economically and operationally prohibitive. However, its role will gradually and predictably shift. From being the primary standard for new integrations, it will become the de facto standard for legacy systems, while new applications will adopt FHIR almost exclusively. Integration engines will become increasingly sophisticated universal translators, allowing both worlds to coexist without apparent friction for clinical users.

Developers who can navigate this technical duality—who understand both the grammar of HL7 v2 and the principles of FHIR—will be the most valuable professionals in a job market that desperately needs to build robust bridges between the past and the future of digital health. It is not about choosing a side in a standards war but about mastering the entire ecosystem to be able to offer solutions that work in the real world, not just in development labs.

7.3 DICOM: MEDICAL IMAGES AND BEYOND

“A medical image without context is like a diagnosis without a medical history: technically present, but clinically useless.”

Imagine a radiologist sitting at his workstation, focused on interpreting a chest CT scan. The image he sees on the screen is far from a simple digital photograph like the ones we take with our phones. It is a complex web of multidimensional information: precise spatial coordinates on three axes, quantified X-ray attenuation values measured in Hounsfield units, time sequences in the case of a dynamic contrast-enhanced study, and an invisible yet crucial layer of metadata describing the patient, the equipment used, the acquisition protocol followed, and the specific technical conditions. This seemingly simple image lies behind a sophisticated standard that enables the CT scanner, the mass storage system, the specialized workstation, and the electronic health record to share not only the pixels but also all the context necessary to interpret them correctly from a medical standpoint. This standard is called DICOM, and it constitutes, after HL7 v2, the most widespread and fundamental protocol in the global healthcare ecosystem.

DICOM, an acronym for Digital Imaging and Communications in Medicine, represents the de facto international standard for comprehensive management, long-term storage, specialized printing, and securtion of medical images in clinical settings.⁽³⁾ Developed and actively maintained by a joint committee that includes the Radiological Society of North America and the American College of Radiology, DICOM transcends the mere definition of a static file format. It is, in fact, a complete and living protocol that rigorously defines how imaging devices must communicate with one another, how images must be structured and stored to ensure their temporal integrity, how associated metadata must be organized to preserve the clinical context, and how the semantic consistency of the information must be ensured throughout its entire lifecycle. Any medical imaging device that aspires to interoperability—whether a simple digital X-ray, a sophisticated multislice computed tomography scanner, a high-field magnetic resonance imaging system, a Doppler ultrasound, or a digital mammography system—inevitably speaks the DICOM language.

7.3.1 THE PHILOSOPHY OF IMAGE-DATA INTEGRATION

What makes DICOM truly special and what radically sets it apart from conventional image formats such as JPEG or PNG is its ability to integrate both the image itself and all the contextual data necessary for its valid medical interpretation into a single standard container. A DICOM file contains not only the pixel matrix representing the visualized anatomy but also an extensive, structured header with complete patient demographic information, comprehensive technical data on the acquisition process, specific details of the equipment used—including calibrations and hardware parameters—the applied algorithmic reconstruction settings, and even semantic relationships with other studies of the same patient or from the same session. This intrinsic integration is not a technical luxury but an absolute necessity for safe and effective clinical practice.

A medical image stripped of its technical and clinical context is practically useless and even potentially dangerous if misinterpreted. Knowing that a certain pulmonary opacity appears in a CT image has no real diagnostic value if we do not know whether the study was performed

with or without intravenous contrast, at what specific phase of the respiratory cycle the slice was acquired to assess mobility, which reconstruction algorithm and which display window were used, or whether the patient was in the supine or lateral decubitus position. DICOM ensures that all this critical information travels inseparably attached to the image so that any system receiving it can accurately reconstruct the acquisition conditions and present it to the specialist with all the information necessary for an informed interpretation.

The internal structure of a DICOM file is organized hierarchically into functional modules that group semantically related attributes.

Its modules are as follows:

- Patient identification module: name, unique identifiers, date of birth, sex, and other essential demographic data.
- Study module: unique study identifier, date and time of acquisition, type of test requested, and referring physician.
- Series module: specific technical conditions for acquiring a related group of images.
- Image module: pixels and their technical encoding characteristics.

This modular organization is not arbitrary; it allows different specialized applications to efficiently access the information they need without having to process and parse the entire file.

An advanced visualization system can prioritize loading the pixels and the image module to quickly display the anatomy to the radiologist, while a file management system can index and search exclusively within the identifying metadata without needing to open and decode the full images, which are large. This functional separation optimizes performance in environments where millions of studies are managed and where different users have distinct access needs. The clinician needs to view the image; the system administrator needs to manage the metadata for archiving and retrieval.

One of the most fundamental and characteristic concepts in the DICOM universe is that of acquisition “modality” or “medium.” Each type of medical imaging equipment is assigned its own modality, identified by a standardized three-letter code that allows for the immediate identification of the technical nature of the study. CR identifies traditional computed radiography, DX refers to modern direct digital radiography, CT designates computed tomography, MR corresponds to nuclear magnetic resonance, US identifies ultrasound or sonography, NM refers to nuclear medicine, and so forth for dozens of specialized techniques. This coding is not merely administrative; each modality has its own defined mandatory and optional specific attributes, reflecting the physical and technical characteristics inherent to each imaging method.

For example, a computed tomography image necessarily includes detailed information on the dose of ionizing radiation administered to the patient, expressed in multiple standardized metrics for dosimetric monitoring, whereas a magnetic resonance image includes complex technical parameters such as the pulse sequence used, repetition and echo times, magnetic field strength, and the characteristics of the receiving coils. This modality-specific semantic richness is precisely what makes DICOM so powerful for specialized clinical practice, but it also explains its considerable technical complexity and the learning curve it presents for developers new to the field.

7.3.2 SERVICE ARCHITECTURE AND PACS ECOSYSTEM

Communication between DICOM devices in a hospital network does not occur via simple file transfer but rather through a sophisticated system of standardized services that define which operations a system can perform on medical objects managed by another. The most fundamental and widely used services are the storage service, which allows an acquisition device to send images to a file server; the query service, which allows a viewer or workstation to ask what studies are available for a specific patient; and the retrieval service, through which the viewer requests and obtains the specific images from a particular study for display. These three services form the basic triangle of DICOM communication.

Beyond these essential services, the standard defines advanced capabilities for printing on film or high-quality diagnostic paper, for managing worklists—that allow radiology technicians to view pending requests directly on the device console, for performing complex procedures involving multiple steps, and for real-time communication of events and alerts between systems. This architecture of modular, well-defined services allows equipment from different manufacturers, designed in different decades and installed in different countries with varying regulations, to interoperate transparently and predictably, provided that both correctly implement the DICOM

conformance profiles.

The central system and nerve center of any modern medical imaging infrastructure is the PACS, an acronym for the Picture Archiving and Communication System. A PACS is an integrated software-hardware platform that receives, temporarily and permanently stores, efficiently distributes, and enables the visualization of medical images exclusively in DICOM format.(4) When a CT scanner completes the acquisition of a study, it automatically sends it to the institutional PACS using the DICOM storage service. From the PACS, any authorized viewing station on the hospital network—or even at secure remote locations—can browse the catalog of available studies and retrieve the specific images needed for diagnostic or clinical work.

PACS also assumes the critical responsibility of long-term storage management, implementing sophisticated data migration policies that automatically move images between high-speed, high-cost disks for frequently accessed recent studies and more economical but slower cold storage systems for older studies that are rarely consulted but must be retained for legal and clinical reasons for years or decades. Without a functional and properly scaled PACS, managing the massive volume of imaging data generated daily in a modern hospital would simply be impossible, and digital radiology as we know it could not exist.

The seamless integration of medical images with electronic health records is another crucial aspect of contemporary healthcare practice. Radiologists do not work in an isolated vacuum; they must have access to the patient's complete medical history, including known allergies—particularly to iodinated contrast agents—previous diagnoses that provide context for current findings and the results of other diagnostic tests. Conversely, clinicians who order imaging tests need to view radiological results coherently integrated into the patient's overall context, not in a separate, disconnected application. For this reason, modern health information systems integrate the DICOM viewer directly into the electronic health record interface.

On the same workspace screen, the physician can simultaneously view the structured radiology report and the images supporting it, navigate through the study's three-dimensional slices, perform basic measurements, and compare the results with those of previous studies of the same patient. This seemingly simple integration requires complex technical orchestration where the PACS and the EHR communicate continuously, typically using HL7 v2 or FHIR for the exchange of demographic and admission data and DICOM for the management of the images themselves. The user perceives a unified experience, but behind the scenes, there are two distinct technological worlds that must be perfectly synchronized.

7.3.3 MODERNIZATION: DICOMWEB AND ARTIFICIAL INTELLIGENCE

One of the most recent, promising, and transformative developments in the DICOM world is the DICOMweb specification ((5)), which represents a modern extension of the traditional standard () designed specifically for the web and cloud era. DICOMweb allows access to medical images and their associated metadata using the same RESTful APIs and exchange formats—such as JSON—that dominate contemporary application development. Traditionally, programmatic access to DICOM images requires specialized low-level libraries, deep knowledge of the original binary protocol, and considerable implementation complexity that limits the ability to develop medical applications.

With DICOMweb, any web developer with experience in standard technologies can build medical applications that access MR or CT images using exactly the same techniques and design patterns they would use to access any other resource on the internet. This democratization of access to medical images is enabling an explosion of innovation, ranging from fully browser-based viewers that require no plugins to sophisticated artificial intelligence tools capable of analyzing millions of images at a massive scale on cloud infrastructures. Specific DICOMweb services such as WADO-RS for object retrieval, QIDO-RS for querying, and STOW-RS for storage map traditional DICOM capabilities to modern web paradigms.(5)

Artificial intelligence and deep learning are radically and rapidly transforming the world of medical imaging, and the DICOM standard is actively evolving to adapt to this new reality. Modern deep learning algorithms can detect potentially malignant pulmonary nodules with sensitivity comparable to or exceeding that of humans, quantify coronary stenosis in angiograms with millimeter precision, automatically segment brain tumors for surgical planning, and perform countless other analytical tasks that previously required exclusively the expert interpretation of a radiologist. These algorithms require massive access to images for training and operation, but they also need to return their results in a structured format that can be seamlessly integrated into

established clinical workflows.

DICOM is progressively incorporating new specialized data structures to represent algorithmically generated anatomical segmentations, structured annotations of findings, derived quantitative measurements, and inference results from artificial intelligence models. These specialized DICOM objects allow AI findings to be permanently stored alongside the original images, displayed as overlays for radiologists on their standard workstations, and included in the patient's legal medical record with full traceability. AI ceases to be an external black box and becomes a first-class citizen of the DICOM ecosystem.

For software developers working specifically with medical images, DICOM undoubtedly presents a steep learning curve, but it is remarkably well documented and supported by a rich ecosystem of open-source tools. There are mature and actively maintained libraries such as DCMTK for C++, fo-dicom for the .NET environment, or pydicom for Python, which greatly facilitate the reading, writing, modification, and manipulation of DICOM files in virtually any modern programming language. There are also complete, open-source DICOM servers such as Orthanc or DCM4CHEE that provide reference implementations of the full protocol.⁽⁶⁾

These servers allow developers to focus on the specific logic of their innovative applications without having to implement the complex state machine and DICOM communication protocol stack from scratch. The quality of the processed images and the absolute accuracy of the associated metadata are critical aspects where no compromises are allowed in DICOM. A seemingly minor error in the spatial orientation of the image, a metadata tag misinterpreted by software, or a loss of clinical information during a compression process can have serious diagnostic consequences and even put the patient's life at risk if they lead to misinterpretation.

For this reason, DICOM systems include robust integrity verification mechanisms using checksums, and exhaustive conformance testing is an absolutely essential and nonnegotiable part of deploying any medical imaging system. The IHE initiative, which we already know for its work with HL7, also defines specific integration profiles for the DICOM ecosystem that specify how systems must behave in specific clinical scenarios such as mammogram reading or worklist management and offers rigorous testing tools to verify that systems technically comply with the specifications.

For physicians who use these systems on a daily basis, DICOM is invisible yet absolutely essential, such as oxygen for breathing. When a radiologist opens a complex study and sees the multiple image series perfectly organized chronologically, with patient data correctly displayed in the headers and with window and level manipulation tools working instantly, they are seeing the tangible result of decades of international standardization work. When a clinician accesses their patient's images from medical records and can navigate smoothly through the thin slices of a three-dimensional CT scan, they are utilizing a complex technological infrastructure that is only possible because of the existence and maturity of DICOM.

The standard's operational transparency—the fact that it simply works without the user having to think about it—is perhaps the best proof of its success and maturity. The future of DICOM is intrinsically linked to the very future of medical imaging as a discipline. Emerging technologies such as dual-energy computed tomography for tissue characterization, functional magnetic resonance imaging for brain mapping, high-resolution molecular imaging, or hybrid techniques such as PET–CT for simultaneous metabolic–anatomical diagnosis all find their reflection and support in new, carefully designed extensions of the standard. Advanced artificial intelligence, immersive three-dimensional visualization, and augmented reality for surgical planning depend fundamentally on DICOM to access the data they need to process and enrich.

Far from being a mature standard in the sense of being static or complete, DICOM continues to actively evolve to adapt to the new clinical and technical needs that constantly arise. In conclusion, it is worth remembering that DICOM is not merely an abstract technical specification but also a fundamental enabler of modern medicine as we practice it. Without DICOM, comprehensive digital radiology could not have developed, teleradiology for specialized support in remote areas would be practically impossible, and integrated multimodal imaging would remain an unfulfilled dream.

Every time a patient does not undergo an unnecessary repeat test because their images from another facility are accessible and viewable owing to standardization, every time a specialist can urgently consult a study from home after hours to make a critical decision, every time an artificial intelligence algorithm helps detect cancer in its early stages when it is still curable, we are witnessing the tangible fruits of this standard that has made universal interoperability in medical imaging possible. DICOM is, in this profound sense, one of those invisible infrastructure technologies that make everything else possible—the very thing we take for granted until it fails.

7.4 SMART ON FHIR: APPS THAT INTEGRATE WITH EPIC AND CERNER

“The best health app in the world is useless if the doctor has to step out of their workflow to use it; seamless integration is the true battleground of clinical innovation.”

Imagine a primary care physician who uses an electronic health records system from one of the major global providers, such as Epic or Cerner, on a daily basis. Dozens of essential features coexist on their screen: appointment scheduling, electronic prescribing, viewing lab results, writing clinical reports, and communicating with other professionals. One day, during a medical conference, they discovered an application developed by an innovative external team that calculates their patients' cardiovascular risk using a validated, novel algorithm that considers genetic factors and emerging biomarkers. To use this potentially useful tool, the doctor would have to completely step out of their usual work environment, open a separate browser window, manually enter all the demographic and clinical data for the patient they are seeing, wait for the calculation, note the result mentally or on paper, and then return to their main system to continue the consultation. This process is tedious, prone to transcription errors, consumes valuable time that cannot be devoted to the patient, and brutally disrupts the cognitive workflow that the physician has developed over the years. Most likely, after trying it a couple of times with initial enthusiasm, the physician will abandon its use and return to traditional methods—not because the application is bad but because the friction of use outweighs its potential value.

This scenario, which has played out countless times over the past decade in the history of digital health, illustrates a fundamental problem that developers and entrepreneurs in the sector have ignored for far too long: there is no point in having the best medical apps, the most accurate algorithms, or the most beautiful interfaces if they do not integrate seamlessly and naturally into the environment where healthcare professionals already work and spend most of their day. The barrier to adoption lies not in the quality of the software but in its operational ubiquity. SMART on the FHIR was created precisely to solve this integration problem in an elegant, standardized, and technically robust way. SMART, which stands for Substitutable Medical Applications, Reusable Technologies, is an open standard initially developed by researchers at Harvard and Boston Children's Hospital that allows applications developed by independent third parties to integrate securely, contextually, and seamlessly within leading electronic health record systems.⁽⁷⁾

Combined with the FHIR for structured access to clinical data, SMART provides a comprehensive framework that allows any developer—from a two-person startup to a large technology corporation—to create applications that can run directly withi, Epic, Cerner, Meditech, or any other system supporting the standard, behaving as if they were native features of the underlying system itself. The clinician can launch the application from a simple button or context menu within their usual interface, and the application automatically receives the full context of the patient they are currently viewing, including essential demographic and clinical identifiers, without requiring the clinician to manually enter any data. The transition between the main system and the integrated application is so seamless that the user perceives a unified experience, not two separate systems that require a mental shift in context.

7.4.1 THE FOUNDATIONAL PRINCIPLES OF THE ECOSYSTEM

The design philosophy of SMART on FHIRs is built on several key principles that together define a vision of open architecture for digital health. The first fundamental principle is substitutability: applications are explicitly designed to be replaceable by others that perform similar or equivalent functions, without the need to modify the underlying EHR system or reconfigure complex point-to-point integrations. This fosters genuine competition on the basis of clinical quality and usability and accelerates innovation, as EHR vendors can offer their customers a broad and diverse ecosystem of specialized applications without having to bear the cost and risk of developing them all in-house. A hospital can switch its cardiovascular risk calculation application to an improved one without having to touch its underlying Epic or Cerner infrastructure.

The second principle is cross-platform reuse: a single application, developed once according to SMART on FHIR standards, can run seamlessly across different EHR systems, provided that they correctly implement the standard. This drastically reduces development costs by eliminating the need to create platform-specific versions and allows the best applications—regardless of their geographic or corporate origin—to reach more clinical users worldwide. A Spanish developer can create a tool that works perfectly in U.S. hospitals using Epic or French hospitals using Cerner

without modifying a single line of code for each market.

The third principle is security by design: access to sensitive patient data is strictly controlled through modern authorization mechanisms based on OAuth2, derived from best practices in the web industry, and the patient can have transparent visibility into which applications are accessing their medical information and for what specific purpose. This security approach is not an afterthought but is built into the architecture from the ground up, recognizing that trust is the most valuable asset in health data exchange.

The technical architecture of the SMART on the FHIR combines several specialized components that work in coordination to deliver this integrated experience. At the heart of the system is the electronic health record system's FHIR server, which exposes structured clinical data via a modern RESTful API, accessible through standard HTTP calls. In addition to this FHIR database, a dedicated SMART authorization service is implemented that centrally manages which applications can access which data, under what contextual conditions, and for how long. The applications themselves are standard web applications, typically developed using technologies such as HTML5, and modern frameworks, which run in the user's browser or in a secure sandboxed container within the EHR system.

When the physician decides to launch an application from their usual interface, an OAuth2 authorization flow is automatically initiated, which, following authentication and with the user's explicit consent, grants the application a temporary access token limited specifically to the current context: the particular patient the physician is seeing at that moment, for the duration of the consultation or the configured session time. This token does not allow the application to access other patients or unrelated historical data, thereby limiting the risk in the event of a compromise. Once authorized, the application can make calls to the FHIR API to obtain the necessary data, process it according to its specific logic, and present results or collect additional input from the physician, all within the same window or panel integrated into the main interface.

One of the most important and distinctive concepts in SMART on FHIR is that of "launch" or "launch context". When an application is launched from the host EHR system, it receives not only the technical access token but also a rich set of information about the current operational context: which specific patient the clinician is currently viewing, the authenticated user's identifier and professional role, the specific visit or care episode they are in, and even the location or department. This contextual information allows the application to configure itself automatically and intelligently, immediately displaying the correct patient data, applying known user preferences, and positioning itself at the appropriate point in the clinical workflow without requiring the physician to make any manual selections or enter identifiers.

The result is a remarkably seamless user experience, where the external application feels like a natural and organic extension of the core system, not an external intrusion. The clinician does not experience the sensation of "switching programs" that generate so much resistance to the adoption of new tools; they simply access additional functionality that magically appears contextualized with the information they are already viewing. This architectural transparency is precisely what makes the clinical adoption of external innovations viable in an environment where the clinician's time and attention are scarce and precious resources.

7.4.2 THE APPLICATION ECOSYSTEM AND THE DEVELOPER EXPERIENCE

The practical and successful use cases of SMART on FHIRs are extraordinarily diverse and have grown exponentially since the standard matured. There are sophisticated clinical decision support applications that calculate specific disease risks, suggest treatments on the basis of updated clinical practice guidelines, or provide real-time alerts about potentially dangerous drug interactions that the base system does not detect. There are specialized applications for the comprehensive management of chronic diseases that help monitor the longitudinal progression of patients with diabetes or hypertension, visualizing trends in clinical values and suggesting preventive interventions based on predictive algorithms. There are applications aimed at clinical research that allow for the identification and recruitment of patients potentially eligible for clinical trials directly from the daily clinical setting without interrupting the flow of care.

Applications for patient education are also proliferating, automatically generating materials tailored to a patient's specific condition, reading comprehension level, and language preferences or advanced visualization tools that transform complex laboratory or imaging data into intuitive graphical representations for the clinic. The list is practically endless and grows every day as more developers discover the platform's possibilities and the relative ease with which they can bring

their innovations into the real clinical setting, without the commercial and technical barriers that previously impeded market access.

For software developers, SMART on the FHIR drastically reduces the historical barriers to entry into the digital health market, which for years was dominated by a few large vendors with closed APIs and opaque integration processes. Before this standard existed, creating an application that integrated reasonably well with Epic or Cerner required complex commercial negotiations that could last years, costly specific partnership agreements, and painful technical adaptations to each vendor's proprietary APIs and incomplete documentation. Each integration was a custom consulting project, not a scalable product. With SMART on the FHIR, any developer with standard modern web development skills can create an application that works on any system supporting the standard, using tools and frameworks they already know.

Detailed implementation guides are publicly available at no cost; there are open-source libraries in multiple languages that facilitate OAuth2 authentication and access to FHIR data, and there are public test environments such as the SMART Launcher where developers can validate their applications with realistic synthetic data before attempting to deploy them in hospital production environments. This democratization of access to the ecosystem has sparked an explosion of innovation from diverse sources: agile startups, hospital IT departments, academic research groups, and even individual developers with specific ideas for particular clinical niches.

Security and privacy are absolutely central to the design of the SMART on the FHIR, not secondary considerations. Applications never have direct access to the underlying electronic health record database, not even in read-only mode; they can only access data through the controlled FHIR API under the constant supervision of the SMART authorization service. Permissions are granted via scope profiles that define, with granular detail, which FHIR resources the application can read or write and in which specific patient or user contexts. The patient can be transparently informed about which applications have access to their data and for what stated purposes and can revoke that access at any time through consent management interfaces.

All access by applications is automatically recorded in immutable audit logs that enable the detection of inappropriate use, anomalous patterns, or violations of institutional policies. This security model, derived from and adapted to best practices in the modern web industry and standards such as OAuth2 and OpenID Connect, provides an adequate and auditable level of protection for sensitive health data, complying with regulations such as HIPAA in the United States or GDPR in Europe, although it is always the responsibility of the end implementer to ensure specific compliance with each jurisdiction.

For large EHR systems such as Epic and Cerner, adopting and promoting SMART on FHIR has been a strategic decision that balances competition and collaboration. It allows them to offer their institutional clients a much richer, more diverse, and up-to-date ecosystem of applications than they could develop internally with their own resources, without having to assume the cost of development, maintenance, and legal liability for each individual application. It also allows them to maintain their position as a central and indispensable platform within the hospital's architecture, since the applications run within their technological environment and rely fundamentally on their APIs to function, thus creating a symbiotic relationship where the success of external applications reinforces the value of the base system.

Major vendors have created their own certified app stores, such as Epic's App Orchard or Cerner's Code, where customers can discover, evaluate, and install SMART on FHIR applications that have passed security reviews, technical compliance checks, and basic clinical validation. These stores operate similarly to smartphone app stores but with more rigorous validation standards given the risks inherent in medical applications. For developers, having a presence in these stores represents a massive and highly credible distribution channel.

7.4.3 ARTIFICIAL INTELLIGENCE AND THE FUTURE OF CLINICAL INTEGRATION

The relationship between SMART on FHIRs and artificial intelligence applied to healthcare is particularly promising and is generating some of the most sophisticated applications in the ecosystem. Machine learning and deep learning models can be easily deployed as SMTs in FHIR applications that, launched directly from the electronic health records, analyze structured patient data in real time and return probabilistic predictions, management recommendations, or risk alerts that are immediately integrated into the clinical decision-making workflow. A radiologist could use an application that employs computer vision to detect and quantify suspicious pulmonary nodules while viewing CT images, receiving an overlaid annotation without leaving

their usual viewer.

A cardiologist could use an application that estimates the risk of major cardiovascular events over the next five years on the basis of the patient's complete history, recent biomarkers, and lifestyle factors, presenting the result as an interpretable probability with highlighted contributing factors. A primary care physician could use an application that suggests specific follow-up tests according to the most up-to-date clinical guidelines, automatically checking for contraindications and prioritizing recommendations on the basis of the patient's profile. All of this occurs without the professional having to leave their usual work environment, with the patient's context already loaded and available, reducing the friction of adopting these emerging technologies to a minimum.

For physicians using these applications in their daily practice, the ideal experience should be completely transparent and technologically invisible. A discreet yet accessible button on the main screen, a context menu that appears at the right moment, an additional tab in the patient interface—and the application is there, offering its specific functionality without perceptible friction or steep learning curves. What happens behind this seemingly simple interface—the complex dance of OAuth2 authorizations, calls to REST APIs, secure data transfer, and algorithmic processing—must remain completely invisible to the clinician. When a SMART on an FHIR application works well, the clinician does not even consciously perceive it as something separate or external; it is simply another feature of their system, as natural as any other feature they use routinely.

This is the design ideal to which SMART on FHIRs aspires: that technological integration be so natural and seamless that it disappears from the user's awareness, and the only thing that remains is tangible clinical utility and improved efficiency or quality of care. The technology becomes transparent, and what remains is pure medical value.

Application governance and curation are critical aspects that hospitals and healthcare systems must carefully consider before opening their doors to this ecosystem. Not all SMTs in FHIR applications are equally reliable, clinically validated, or secure, and the use of an open standard does not guarantee their quality, efficacy, or safety for patients. Hospitals need to establish clear institutional processes to evaluate applications before allowing their productive use, verifying who developed them and their reputation in the market, what specific data they access and whether that access is proportionate to the stated purpose, what guarantees of technical support and updates they offer, and how security vulnerabilities discovered later are managed.

Some electronic health record (EHR) system vendors offer certification programs that help hospitals select prevalidated, trusted applications, but the ultimate responsibility for patient safety and regulatory compliance always rests with the institution authorizing the deployment. Governance is not an obstacle to innovation but rather a necessary framework for ensuring that innovation is sustainable and safe.

The evolutionary future of SMART on FHIRs is intrinsically linked to the future of FHIRs as a data standard and to that of EHR systems as open platforms. As more vendors—including those outside the U.S. market dominated by Epic and Cerner—adopt the standard and as more advanced functionalities are exposed through increasingly rich FHIR APIs, the application ecosystem will continue to grow in complexity and sophistication. We are likely to see the emergence of entire categories of applications that we cannot even fully imagine today, leveraging the convergence of massive amounts of clinical data, advanced artificial intelligence, applied genomics, and integrated user experience. SMART on the FHIR is not just a technical specification; it is a vision of how innovation in digital health can be democratized and distributed, allowing the world's best developers, regardless of their location or size, to contribute to improving global healthcare.

SMART on the FHIR, while powerful, does not magically solve all healthcare integration problems. It remains absolutely necessary for EHR systems to expose the right, high-quality data through their APIs; for FHIR profiles to be welldefined and consistently implemented; for authorization services to function correctly without creating bottlenecks; and for the applications themselves to be clinically valid, usable, and maintained. However, SMART on FHIR provides the essential architectural framework that makes everything else possible—the technical and semantic scaffolding upon which to build an open, competitive, and secure innovation ecosystem. In a world where medical technology is advancing ever faster and where the need for specialized solutions exceeds the capacity of any individual provider, having a standard that allows new solutions to be seamlessly and securely integrated with existing ones is likely the most valuable architectural contribution we can make to the future of digital health.

If SMART on the FHIR has succeeded in opening up electronic health record systems to external innovation, how should we redefine models of clinical validation and legal liability when

these third- r applications make decisions that directly affect patient safety, and who bears the blame when an integrated algorithm makes a diagnostic error within an established provider's environment?

7.5 APIS FOR WEARABLE DEVICES AND CONNECTED DEVICES

“The watch on the patient’s wrist generates more data in a single day than their entire previous medical history; the challenge is not a lack of information, but transforming that torrent into actionable medical knowledge.”

Imagine a patient with chronic heart failure who has just been discharged from the hospital after being treated for acute decompensation. On their wrists, they now wear a state-of-the-art smartwatch that continuously monitors not only their heart rate and physical activity level but also, using advanced bioimpedance and signal analysis algorithms, estimates subtle changes in their total body fluid volume—one of the earliest indicators of cardiac deterioration. In the bathroom of their home, a connected smart scale automatically records their body weight every morning before breakfast, transmitting the readings to the cloud. On their night stand, an automatic wrist blood pressure monitor takes scheduled readings of blood pressure and resting heart rate. All this information, which is continuously generated in the patient’s natural environment, could be extraordinarily valuable to their primary care physician and cardiologist, allowing them to detect early signs of cardiac decompensation days before the patient notices subjective symptoms and thus avoiding costly and potentially dangerous hospital readmission. However, for this vision of predictive medicine to become a reality, the data generated by these heterogeneous devices must traverse multiple technological layers, ultimately reaching the institutional electronic health record system, integrate coherently with the rest of the patient’s information, and be presented in a way that allows the physician to interpret it quickly and act on it with confidence. This is exactly what wearables and connected device APIs are about: building technical and semantic bridges between the consumer ecosystem of personal devices and professional healthcare infrastructure.

The massive explosion of connected devices in the health and wellness sector represents one of the most transformative technological phenomena of the past decade in medicine. It is estimated that hundreds of millions of active wearables are already in use globally, ranging from multigenerational smartwatches to adhesive biosensor skin patches, continuous ambulatory glucose monitors, smart inhalers with adherence sensors, and long-term implantable cardiac monitoring devices. All of these factors generate a continuous, voluminous, and multidimensional flow of data that, when properly processed, contextualized, and integrated, can revolutionize population-based preventive medicine, the proactive management of complex chronic diseases, and the conduct of large-scale clinical research using real-world data. However, this tsunami of personal data also poses formidable challenges regarding technical interoperability, information security, clinical validation of algorithms, and ethical privacy management on a scale that traditional healthcare systems have never faced before.

7.5.1 THE TOWER OF BABEL OF MANUFACTURERS AND THE EMERGENCE OF STANDARDS

The first unavoidable issue that arises when seriously addressing the integration of wearables is the overwhelming diversity of manufacturers, hardware architectures, and proprietary communication protocols. Through its Fitbit and Android ecosystem, Samsung, Garmin, Withings, Oura, and dozens of other specialized manufacturers all have their own devices with specific sensors, their own companion mobile apps with distinctive user interfaces, and their own idiosyncratic ways of exposing data to external developers. For a developer aiming to build a clinical application that universally collects information from all of them, the initial task can seem daunting, requiring the mastery of multiple SDKs, different authentication models, and the mapping of incompatible data formats. Fortunately, in recent years, standardization initiatives and intermediary platforms have emerged that progressively facilitate this integration, reducing the fragmentation that threatens to paralyze the ecosystem.

The most prominent platforms in this space are Apple’s HealthKit, which provides a centralized repository of health data on iOS devices, and Google Fit, which offers similar functionality for the Android ecosystem. These operating system APIs allow authorized applications to read and write a wide variety of standardized data types, such as step count, resting and active heart rate, sleep patterns and quality, body weight, blood glucose levels, blood pressure, oxygen saturation,

and many other biometric parameters. These APIs follow modern authorization models based on granular permissions, where the user maintains explicit control over which specific apps can access which categories of sensitive data and for how long and can revoke these permissions at any time from a centralized privacy interface.

The next critical step in the integration chain is to transport these data from consumer manufacturers' applications to institutional electronic health records systems. This is where the FHIR standard, which we have explored in the previous sections, plays an absolutely central role as the lingua franca of digital health. An intermediary application can read data from the wearable via the manufacturer's API or the operating system's health platform, semantically transform it into standardized FHIR resources (observation for point-in-time or time-series biometric readings, questionnaire response for health questionnaire responses, device for information on the hardware used, etc.), and finally send it to the hospital's FHIR server using precisely the SMART on the FHIR APIs we developed in the previous section. The architectural result is that data generated by the patient in the privacy of their home during their daily activities are integrated into their official electronic health records in the same structured and secure manner as laboratory results processed at the hospital or imaging tests performed in radiology.

However, the technical integration does not end there, and in fact, that is only the beginning of the potential value. Once the data from wearables are structured within the electronic health record, they can be leveraged for multiple sophisticated clinical purposes. It can feed remote monitoring dashboards that the physician systematically reviews before a scheduled visit, providing a longitudinal context that enriches the in-person consultation. It can generate intelligent automatic alerts when analysis algorithms detect abnormal patterns or concerning trends, such as sudden weight gain in a cardiac patient, suggesting fluid retention. It can be incorporated into machine learning predictive models that dynamically stratify patient risk and prioritize proactive care. They can even be used to adjust pharmacological treatments remotely—through automated titration protocols—without the need for the patient to physically attend a consultation. The ultimate value of wearable data does not lie in the bits and bytes themselves but in the preventive and personalized clinical actions they enable.

7.5.2 REGULATED MEDICAL DEVICES AND DATA QUALITY

One of the most advanced and clinically mature use cases involves connected medical devices that, unlike general-purpose wearables, already incorporate capabilities for direct and secure communication with professional healthcare systems, designed from the outset for clinical integration. Modern pacemakers with telemetry capabilities can send detailed automatic reports to cardiologists when their internal algorithms detect significant arrhythmias or changes in pacing parameters. Next-generation continuous glucose monitors can share real-time readings not only with the patient but also with the hospital's multidisciplinary diabetes team, enabling immediate insulin adjustments. Smart inhalers with usage sensors can accurately record when, where and how the patient uses them, helping to identify patterns of environmental triggers for asthma attacks and assess actual treatment adherence. These devices are not mass-market wearables but are regulated medical devices with CE marking or FDA clearance, and their APIs are specifically designed to meet the safety, reliability, and traceability requirements of professional clinical settings.

Information security and patient privacy are, as in the entire digital health ecosystem, absolutely critical and nonnegotiable aspects. Even if it comes from consumer devices, the data generated by wearables are unquestionably sensitive personal health data and are therefore subject to the same strict regulatory protections as any other traditional clinical data. This means that all communications between devices, intermediary applications, and hospital systems must be encrypted using robust protocols such as TLS; access must be rigorously authenticated through mechanisms such as OAuth2 and managed via role-based authorization; and there must be immutable audit logs that allow tracking of who accessed which specific data and at what time. It also means that the patient must maintain full control over their data and be able to make informed decisions about what to share, with which institutions or applications, for how long, and for what specific purposes. Wearable manufacturers' platforms typically provide user interfaces for patients to manage these permissions in a granular manner, but the ultimate responsibility for technically respecting these consent decisions lies with the applications that consume them and the systems that store them.

The clinical quality and metrological reliability of wearable data is another aspect that

deserves critical attention and is often underestimated by technology enthusiasts. Not all devices available on the market are equally accurate, validated, or robust, and real-world patient usage conditions can significantly affect the quality of measurements. A smartwatch that measures heart rate using photoplethysmography can produce erroneous readings or artifacts if it is not properly positioned on the wrist, if the patient moves vigorously during the measurement, or if the skin is damp. An activity tracker might count movements that are not actually walking—such as arm swinging while driving—as steps. Clinical systems that receive and interpret these data must be explicitly aware of these technical limitations and present them to healthcare professionals with the appropriate caveats and context to enable informed critical interpretation.

A physician who sees an anomalous reading on a telemonitoring dashboard must be able to easily access the relevant technical context: is this an isolated single value or a sustained trend over several days? Was the device recently calibrated according to the manufacturer's specifications? Are there environmental or usage factors that could explain an apparently erroneous reading? What is the sensor's stated accuracy for the observed range of values? Without this contextualization, the risk of overdiagnosis—of false alarms that cause patient anxiety and strain the system—or of underdiagnosis due to widespread mistrust is considerable.

The integration of data from wearables also poses significant technical challenges in terms of volume, velocity, and variety that traditional health information systems were not designed to handle. A patient with an ambulatory continuous glucose monitor can easily generate hundreds of individual readings per day, each with multiple associated data dimensions. Multiplied by thousands of patients being monitored simultaneously in a region or healthcare system, the cumulative volume of data can reach large data magnitudes that challenge conventional clinical database architectures.

Systems must be capable of ingesting, storing, and processing this massive volume without degrading the operational performance of critical clinical applications, and they must present the resulting information in an aggregated and summarized form that is practically useful to clinicians, who obviously cannot manually review hundreds of individual readings per patient during a fifteen-minute consultation. This requires sophisticated capabilities for temporal aggregation, trend visualization, intelligent alert thresholding, and analysis algorithms that proactively highlight what is relevant and suppress noise.

7.5.3 THE FUTURE OF CONNECTED, PATIENT-CENTERED MEDICINE

For software developers working specifically with wearables and connected devices, the variability of APIs and formats remains a constant operational challenge that requires flexible and well-designed software architectures. Each manufacturer maintains its own way of semantically representing data, its own authentication and authorization mechanisms, its own rate-limiting and access quota restrictions, and its own evolving privacy policies. Fortunately, specialized libraries and aggregation services are emerging in the market that abstract much of this underlying complexity, providing a unified and standardized interface for multiple devices and manufacturers. Industry-level standardization initiatives are also emerging, such as those promoted by the Personal Connected Health Alliance, which seek to define common formats and interoperable protocols for wearable data, facilitating their seamless integration with institutional healthcare systems and reducing the development burden for integrators.

The synergistic relationship between wearables and telemedicine is particularly relevant and transformative for the future of care models. Compared with traditional sporadic in-person visits, connected devices effectively function as the physician's extended eyes, ears, and senses in the patient's home, within their natural living environment, providing continuous and contextualized information that enables much closer, more frequent, and personalized clinical monitoring. This is particularly valuable for patients with complex chronic conditions requiring intensive monitoring; for frail elderly individuals living alone, early detection of functional decline can prevent institutionalization; for patients in remote rural areas with limited geographic access to healthcare facilities; and for populations with reduced mobility. The intelligent combination of wearables, synchronous and asynchronous telemedicine, and clinical decision algorithms promises a more proactive, preventive, personalized, and equitable healthcare system.

For patients who generate these data, the ability to actively share them with their doctors can be both an empowering opportunity and a legitimate source of anxiety and concern. On the one hand, it allows them to feel more involved and engaged in their own healthcare, better monitored and safer, and more connected to their care team between in-person visits. On the

other hand, it can raise valid concerns about the privacy of their most intimate data, about who else might access it without their knowledge, about how it will be used in the long term, and about the risks of discrimination or stigmatization if certain sensitive data are leaked or used out of context. It is the ethical and legal responsibility of developers and healthcare systems to design user experiences that fully respect patient autonomy, are radically transparent about the specific use of data, and offer clear, simple, and effective controls to manage sharing permissions at any time.

The future of wearables and connected health devices is extraordinarily promising and is accelerating continuously. Devices will become progressively more metrologically accurate, smaller and more discreet, and more autonomous in terms of battery life and local processing capacity. They will measure an increasing number of physiological parameters with greater clinical reliability over longer periods without the need for recharging or recalibration. They will integrate more naturally and invisibly with smart clothing, with everyday objects in the connected home, and even with biocompatible subcutaneous implants. They will generate volumes of personal data that seem unimaginable today but that will be the essential raw material for truly personalized, predictive, and preventive medicine. For developers who know how to navigate this complex ecosystem of heterogeneous APIs, emerging standards, healthcare regulations, and ethical considerations, the opportunities for positive clinical impact are immense and are only just beginning to be explored.

Technology, however sophisticated it may be, remains merely a means to a higher end. The ultimate goal of integrating wearables and connected devices with healthcare systems is not simply to accumulate more data for their own sake but also to empower professionals to make better clinical decisions and patients to better manage their health. Developers building these integrations must always maintain an obsessive focus on tangible clinical value—on how the information flowing from patients' devices can concretely help professionals provide better, more personalized, and more proactive care. That is, and must remain, the ethical and practical compass guiding all innovation efforts in this space.

If devices in patients' homes generate potentially critical clinical data that flows continuously into hospital systems, how can we redefine the boundaries of medical liability and professional–patient relationship models when monitoring is constant and when intervention can be remote and even automated?

REFERENCES

- Ayaz M, Pasha MF, Alzahrani MY, Budiarto R, Stiawan D. The Fast Health Interoperability Resources (FHIR) Standard: Systematic Literature Review of Implementations, Applications, Challenges and Opportunities. *JMIR Med Informatics* [Internet]. July 1, 2021 [cited May 13, 2026];9(7):e21929. Available from: <https://pmc.ncbi.nlm.nih.gov/articles/PMC8367140/>
- Simborg DW. Reflections on the history of interoperability in hospitals. *J Am Med Inform Assoc* [Internet]. December 1, 2023 [cited May 13, 2026];30(12):2083. Available from: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10654835/>
- Bidgood WD, Horii SC, Prior FW, Van Syckle DE. Understanding and Using DICOM, the Data Interchange Standard for Biomedical Imaging. *J Am Med Informatics Assoc* [Internet]. 1997 [cited 2026 May 13];4(3):199. Available from: <https://pmc.ncbi.nlm.nih.gov/articles/PMC61235/>
- Andriole KP. Picture archiving and communication systems: past, present, and future. *J Med Imaging* [Internet]. December 28, 2023 [cited May 13, 2026];10(6):061405. Available from: <https://pmc.ncbi.nlm.nih.gov/articles/PMC10754358/>
- Genereaux BW, Dennison DK, Ho K, Horn R, Silver EL, O'Donnell K, et al. DICOMweb™ : Background and Application of the Web Standard for Medical Imaging. *J Digit Imaging* [Internet]. June 1, 2018 [cited May 13, 2026];31(3):321. Available at: <https://pmc.ncbi.nlm.nih.gov/articles/PMC5959831/>
- DCMTK - dicom.offis.de [Internet]. [cited 2026-05-13]. Available from: <https://dicom.offis.de/dcmstk.php.en>
- Mandel JC, Kreda DA, Mandl KD, Kohane IS, Ramoni RB. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. *J Am Med Inform Assoc* [Internet]. September 1, 2016 [cited May 13, 2026];23(5):899-908. Available at: <https://pubmed.ncbi.nlm.nih.gov/26911829/>