

Chapter 9 / Capítulo 9

Primum non nocere: Software engineering
for medicine
ISBN: 978-9915-9876-8-2
DOI: 10.62486/978-9915-9876-8-2.Ch09
Pages: 151-166



© 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 9. Clinical Cybersecurity

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.

“IT security in a hospital is not just another technical problem: it is an extension of clinical safety. A security breach that renders a medical record inaccessible, alters a patient's data, or allows unauthorized access to sensitive information is not an IT incident: it is a clinical incident with potentially serious consequences for those who trusted the system to protect them.”

Imagine that a hospital suffers a ransomware attack. Attackers have managed to encrypt servers containing medical records, laboratory results, diagnostic images, and prescription systems. Doctors cannot access their patients' information. Nurses do not know what medication has been prescribed. The emergency room has to operate using paper and memory, with the risk of errors that this entails. Scheduled surgeries are canceled. Chronic patients lost their follow-up care. This scenario, which a decade ago appeared to be science fiction, has played out dozens of times in hospitals around the world. Cybersecurity has ceased to be a peripheral concern of IT departments and has become a central issue of clinical safety.

The healthcare sector has become one of the prime targets for cyberattackers. Hospitals store extremely valuable data, are willing to pay ransoms to regain access to their systems, and struggle to implement robust security measures because of the complexity of their environments and the age of much of their equipment. Added to this is the unique reality that healthcare systems cannot afford to be down: while a store can close for a day because of a cyberattack, an emergency room cannot. This asymmetry makes hospitals prime targets.

In this chapter, we explore the fundamental aspects of cybersecurity in healthcare environments. We begin with the regulatory framework that protects sensitive data: the HIPAA in the United States, the GDPR in Europe, and local regulations in other countries. We will continue by analyzing why hospitals have become such an attractive and vulnerable attack vector. We address the principles of security by design, which should guide the development of any healthcare system. Finally, we look at the specific mechanisms of authentication, authorization, and auditing that allow us to control who accesses what information and leaves a trail of that activity. The goal is not to turn the reader into a security expert but to equip them with the awareness necessary to understand that security is everyone's responsibility.

9.1 SENSITIVE DATA: HIPAA, GDPR, AND LOCAL REGULATIONS

There is a natural tendency, especially among software developers accustomed to working with commercial or financial data, to believe that health data are essentially similar to any other personal data and are just slightly more sensitive and therefore require slightly more care. This perception is dangerously naive and underestimates the unique nature of health information. A person's medical information is likely the most intimate, personal, and potentially harmful type of data about them. It reveals not only diagnosed medical conditions but also private lifestyle habits, complex family relationships, hereditary genetic predispositions, and profound aspects of personality and psychological functioning that not even the person's closest friends are aware of. A detailed psychiatric history documenting episodes of major depression or anxiety disorders, a diagnosis of HIV infection, a record of therapeutic abortion, a failed assisted fertility treatment, and a sexually transmitted disease: all of these are clinical data that, if disclosed without explicit consent or used for unauthorized purposes, can cause immense, lasting, and sometimes irreparable harm to the affected person, ranging from explicit or veiled workplace discrimination to severe social stigmatization, devastating family conflicts, and profound psychological damage resulting from the violation of the most intimate privacy.

For this fundamental reason of human dignity and social justice, over the course of decades, democratic societies have developed specific, complex, and constantly evolving legal frameworks to rigorously protect health information. The best known internationally, especially in the English-speaking world, is the Health Insurance Portability and Accountability Act (HIPAA), the U.S. Health Insurance Portability and Accountability Act, passed in 1996 and subsequently amended. The HIPAA establishes a comprehensive set of mandatory requirements to protect protected health information (PHI),(1) including detailed privacy rules that strictly limit permitted uses and authorized disclosures of data; technical security rules that require specific administrative, physical, and technical safeguards; and breach notification rules that require covered organizations to promptly notify affected patients and regulatory authorities when a security incident that compromises data occurs. For more than two decades, the HIPAA has been the global benchmark for health data protection, and many countries have adopted structural elements to develop their own national legislation tailored to their context.

In Europe, the General Data Protection Regulation (GDPR) of 2016 established a standard of protection that is even more demanding and comprehensive than the HIPAA in many respects. (2) The GDPR explicitly considers health data, broadly defined, as a “special category” of personal data subject to enhanced protection and additional restrictions. This means that, by legal default, these data cannot be processed without first obtaining explicit, free, specific, informed, and unambiguous consent from the data subject, except for very limited and strictly interpreted exceptions such as the provision of direct healthcare to the data subject themselves or scientific research with prior ethical approval from a competent committee. The GDPR also introduces new and powerful rights for European citizens, such as the right to be forgotten (to have their data erased when it is no longer necessary), the right to data portability (to receive their data in a structured format and transfer it to other controllers), and the right to be informed without an undue delay of security breaches affecting their data, within a maximum of 72 hours of detection. Penalties for noncompliance with the GDPR can reach 20 million euros or 4% of the company's annual global turnover, which has effectively focused the attention of healthcare organizations across Europe on strict compliance.

9.1.1 THE GLOBAL REGULATORY LANDSCAPE AND UNIVERSAL PRINCIPLES

In Latin America, the regulatory landscape varies significantly across countries but shows a clear and growing trend toward progressive harmonization with international protection standards. Argentina was a pioneer in the region with its Personal Data Protection Law No. 25,326 of 2000, which explicitly includes health data as sensitive data subject to enhanced protection. Brazil passed its General Data Protection Law (LGPD) in 2018, structurally inspired by the European GDPR, which also grants special protection and requires specific consent for health data.(3) Mexico has the Federal Law on the Protection of Personal Data Held by Private Parties, supplemented by various state laws that harmonize the framework. Specific legal frameworks for Colombia, Chile, Peru, Uruguay, and other Latin American countries have been developed, with varying degrees of maturity, technical requirements, and capacity for effective enforcement.(4) For healthcare software developers working on transnational projects or within the region, having a detailed understanding of the specific regulations currently in force and under development in each country of operation is absolutely essential, as differences can be significant in areas such as notification deadlines, the legal grounds permitted for processing, or the applicable penalties.

In Spain, the applicable regulations combine the European framework of the GDPR with the Organic Law on Data Protection and Guarantee of Digital Rights (LOPDGDD) 3/2018,(5), which elaborates on specific aspects of the GDPR within the Spanish context and adds relevant provisions such as enhanced protection of health data in the workplace or insurance sectors, where abuses have historically occurred. It also explicitly establishes the obligation to conduct data protection impact assessments (DPIAs) when data processing may pose a high risk to individuals' rights and freedoms, which is practically always the case in healthcare projects involving large-scale health data or emerging technologies such as artificial intelligence. The Spanish Data Protection Agency (AEPD) has issued specific guidelines for the healthcare sector that should be understood and followed, including technical recommendations on anonymization, pseudonymization, and security.

Beyond specific implementation differences across jurisdictions, all these global regulations share common fundamental principles that healthcare software developers must deeply internalize and apply in their daily work. The principles are as follows:

Legitimacy through consent: Given their sensitive nature, health data cannot be used for any purpose without an appropriate legal basis, with the patient's explicit consent being the general rule. This consent must be truly informed (the patient must genuinely understand what their data will be used for and what that entails), specific to each purpose (generic, open-ended consent allowing for any future use is not valid), and revocable at any time (the patient may change their mind and withdraw their consent, with effects for the future but not retroactively). Implementing this in software requires designing robust technical mechanisms to capture consent in an auditable manner, storing it securely and in a retrievable format, strictly respecting it in all processing operations, and allowing for its effective revocation while managing the consequences.

Data minimization: Only data strictly necessary for the specific, stated purpose should be collected, stored, and processed—nothing more. This entails designing systems that do not request or store more information than is essential for the clinical or administrative function they

perform and that implement automatic processes for deletion or anonymization when the data are no longer necessary for that purpose. A system that accumulates patient data indefinitely “just in case” it might be useful in the future, or out of mere technological conservatism, is flagrantly violating this principle and accumulates unnecessary risks.

Technical and organizational security: Data must be protected with specific, proportionate, and demonstrably adequate measures commensurate with the assessed level of risk. This necessarily includes strong encryption of data in transit between systems and at rest in storage, strict access control based on the principle of least privilege, continuous auditing of all significant operations, and the existence of documented security incident response plans. Regulations generally do not prescribe specific universal technical measures but rather require a documented risk analysis to determine the appropriate measures in each specific case, which implies responsibility and technical judgment on the part of the implementers.

Active transparency toward patients: citizens have the right to know, in a clear and accessible manner, what data are held about them, for what specific purposes it is used, who has accessed it and when, and with which third parties it has been shared. This means that health information systems must be technically capable of providing this information in a way that is understandable to nontechnical users and of efficiently handling requests from patients to exercise their rights of access, rectification, erasure, and portability.

Proactive and demonstrable accountability: It is not enough to passively comply with the rules; organizations must actively and documentarily demonstrate that they are complying. This involves maintaining detailed records of processing activities, documenting privacy-related design decisions, conducting periodic impact assessments, and always being prepared to demonstrate to regulators that appropriate safeguards have been implemented. This principle of accountability is central to the GDPR and is changing the culture of healthcare data management.

9.1.2 TECHNICAL CHALLENGES AND SHARED RESPONSIBILITY

For software developers working on healthcare applications, translating these abstract principles into concrete code and architectural decisions is neither trivial nor free of dilemmas. Implementing consent robustly means designing not only user interfaces to capture it but also data structures to store it in an immutable and queryable way, access control mechanisms that dynamically respect the consent status in every operation, and processes to manage revocation along with its implications for ongoing processing. Implementing data minimization means carefully considering which fields are truly necessary when designing database schemas, establishing retention policies with automatic deletion, and resisting the temptation to hoard data “just in case.” Implementing transparency means building patient interfaces that allow them to view and download their data in an understandable way and logging systems that record all access operations with the detail necessary to respond to information requests.

A particularly complex aspect in practice is the granular management of consent in environments where data must be legitimately shared among multiple organizations for comprehensive patient care. A patient may consent to their data being used for scientific research, but only for certain specific types of studies (for example, cancer research but not psychiatric research), and only by certain institutions of recognized prestige and only for a specified period of time. Managing these finely granular consents and ensuring that they are technically respected at every node in an interoperability network is a major technical and organizational challenge that is driving innovation in standards. Technical specifications and emerging technologies specific to this problem are emerging, such as HL7 FHIR consent profiles that allow consent to be represented as an automatically processable structured resource, but their widespread implementation is still in its infancy and requires further development.

Another critical area where developers bear direct responsibility is the reporting of security breaches. When a security incident occurs that affects patients' personal data, regulations require notifying those affected and regulatory authorities within very short timeframes (72 hours in the case of the GDPR from detection, with potential penalties for delays). This means that organizations must have the technical and organizational capabilities to quickly detect security incidents through proactive monitoring, technically investigate their exact scope to determine whether they constitute a reportable breach, and efficiently communicate with those affected using clear and useful information. This is not something that can be improvised when an incident occurs; plans, procedures, communication templates, and response teams must be prepared in advance and regularly tested through drills.

The tense relationship between data protection and scientific research is particularly delicate and requires balance. Accumulated health data serve as a goldmine for medical and epidemiological research, and its secondary use can generate knowledge that benefits society as a whole and future generations of patients. However, this use must be carefully balanced against the right of current patients to control their most intimate information. Regulations establish legitimate avenues for using data for research purposes without explicit individual consent under certain strict conditions (when obtaining consent is impractical because of the size of the cohort, when there is a demonstrated relevant public interest and when effective anonymization is guaranteed). However, these avenues are subject to strict formal requirements for ethical approval, security, and transparency, and their practical application always requires specialized legal advice and case-by-case evaluation.

For clinicians and healthcare professionals involved in projects to develop or acquire information systems, a thorough understanding of these regulations helps them be aware of their own legal and ethical responsibilities. When a physician accesses a patient's electronic health record, they must do so solely in the context of providing direct care to that patient, never out of personal curiosity or to spy on acquaintances. When a researcher requests data for a study, they must ensure that they have all the necessary ethical and legal authorizations before accessing any data. When a manager decides to share data with another institution for patient care, they must verify that appropriate confidentiality agreements and legal bases are in place. Data protection is not just a matter for specialized lawyers or security engineers; it is a shared responsibility of everyone who handles health information in any way, from the physician to the administrative staff to the developer.

These regulations are constantly and rapidly evolving, and developers must proactively stay up to date through ongoing training. The GDPR is subject to periodic revisions and interpretive guidelines that clarify its application to new cases. In the United States, the HIPAA is updated through additional rules that adapt its requirements to emerging technologies. In Europe, the European Health Data Space (EHDS) is being actively developed, which will create a detailed, specific regulatory framework for the cross-border exchange of health data for care and research purposes. In many countries, sector-specific laws are emerging for technologically complex fields such as genomics, telemedicine, or medical artificial intelligence algorithms. Staying up to date in this changing environment is an ongoing professional challenge, but it is absolutely essential for building and operating systems that remain legal, ethical, and socially acceptable throughout their lifecycle.

Health data protection laws exist because there have been serious historical abuses because health information was systematically used to discriminate in employment, to socially stigmatize, and to harm vulnerable people. Complying with the letter of the law is the ethical minimum we can do as professionals; what is desirable and aspirational is to go beyond the legal minimum, building systems that respect patients' privacy and autonomy not because a regulatory standard requires it but because we recognize that it is the morally correct and professionally ethical thing to do. Patients entrust us—as developers, clinicians, and institutions—with their most intimate and vulnerable information; they deserve for us to treat it with the utmost respect, care, and protection possible, today and in the future.

9.2 THE HOSPITAL AS A VECTOR OF ATTACK

Let's imagine a modern, mid-sized hospital in a European city. Hundreds of devices connected to the institutional network coexist daily within its facilities: desktop computers in outpatient clinics, robust servers in the data center, multiparameter monitors in intensive care units, networked programmable infusion pumps, MRI machines with complex control systems, PACS image storage systems, smartphones for medical and nursing staff, tablets for bedside nursing management, and a long list of other medical and administrative technologies. Many of these critical devices were originally designed without considering cybersecurity as a design priority; they run obsolete operating systems that no longer receive security patches and cannot be updated regularly because of manufacturer restrictions or regulatory certifications. The communication network connecting them has grown organically over decades, with successive patches, improvised additions, and unplanned expansions that have created a tangled web of technical and logical dependencies that is difficult to fully audit. Healthcare staff, legitimately overburdened with clinical and administrative work, cannot always follow best practices in terms of IT security, and cyberattackers are well aware of this and systematically exploit it. Far from

being the impregnable digital fortress that patients might imagine, the hospital has become a vulnerable, attractive, and highly profitable target for organized cybercrime.

The **healthcare sector** presents a unique and dangerous combination of structural factors that make it particularly attractive and profitable for all kinds of cybercriminals, from organized ransomware groups to sophisticated state-sponsored actors. The first factor is the high intrinsic value of medical data in black markets. Complete medical records, with all their personal, clinical, and financial information, are sold on underground forums at prices far higher than those of simple financial data such as credit card numbers. With a stolen credit card number, a criminal can make a few fraudulent purchases before the victim detects it and the bank cancels the card, limiting the damage and the criminal's profit. With a complete medical record, however, one can construct a solid false identity for health or life insurance fraud, obtain illegally prescribed controlled substances, commit tax or employment fraud, or even blackmail the affected person over stigmatizing conditions. Health data do not expire nor can they be "changed" by the victim: once it is circulating on the dark web, it can be exploited by multiple actors for years without the victim's knowledge or ability to mitigate damage.

Another factor that makes the healthcare sector irresistible to attackers is the extreme willingness of hospitals to pay for ransoms when they are compromised. Owing to the critical nature of their mission, hospitals simply cannot afford to be offline for extended periods without potentially lethal consequences for patients. When a ransomware attack encrypts critical systems for medical records, surgery scheduling, or lab results, hospital management faces a terrible, high-stakes ethical decision: pay the ransom demanded in cryptocurrency and regain access quickly or refuse on principle and face the consequences of days or even weeks without essential IT systems for patient care. Historically, many institutions have chosen to pay, even though they know this fuels the cybercrime business and attracts more attackers to the sector, and even paying does not technically guarantee that the data will be recovered or that it will not be published anyway. The pressure to provide care is so great that many organizations succumb to extortion.

A medium-sized hospital can also easily have hundreds of different software applications, from dozens of different vendors, with different security architectures, different update and patching cycles, different licensing and support models, and different contractual parties. Managing this technological heterogeneity from a security perspective is enormously complex and costly, and any weak link in this long chain can be exploited by attackers as an initial entry point. An MRI machine purchased fifteen years ago, running embedded Windows XP and no longer receiving security updates from the manufacturer but that remains connected to the hospital network because it needs to share images with the PACS, can be the unnoticed entry point that compromises the entire corporate system.

Many certified medical devices have regulated lifespans of ten, fifteen, or even twenty years and were designed at a time when cybersecurity was neither a design priority nor a regulatory requirement. They often run outdated versions of Windows, Linux, or proprietary operating systems that no longer receive security support from the original software manufacturers. Medical device manufacturers do not always provide security updates throughout the entire lifespan, and hospitals cannot afford or clinically justify replacing equipment that functions perfectly from a medical standpoint solely for cybersecurity reasons. The result is a fleet of inherently vulnerable connected medical devices that attackers know well, actively map, and routinely exploit.

The prevailing organizational culture in many hospitals, where cybersecurity is often perceived as a bureaucratic obstacle to urgent clinical work, rather than as essential protection for patients. Complex password policies are ignored or shared with "streamline" access; privileged access is not revoked in a timely manner when staff leave or change roles; new devices are connected to the network without going through established security checks; and USB storage devices are exchanged among staff without precautions or scanning. This culture—understandable in organizations where the stated top priority is caring for sick patients and where time is literally life—creates systemic vulnerabilities that attackers exploit on a massive scale through social engineering. Changing this culture requires a sustained effort in awareness-raising, ongoing training, and, crucially, designing security systems that are as unobtrusive as possible in critical clinical workflows.

9.2.1 THE EXPANDING ATTACK SURFACE AND CLINICAL CONSEQUENCES

Telemedicine for remote consultations, connected medical devices for home monitoring, remote access for professionals from their homes, and mandatory integration with other centers

to ensure continuity of care—all these legitimate and desirable trends exponentially increase the hospital's total attack surface. Every new authorized external connection is a potential entry point for attackers. These connections are often established under the pressure of immediate clinical need, without adequate prior analysis of the security risks they introduce, simply because the demand for care is urgently needed. The result is an increasingly porous and und hospital network, with multiple weak points that attackers can systematically explore and exploit.

In far too many institutions, all devices logically share the same flat network such that an attacker who initially gains entry through a weak point can move laterally without significant obstacles until it reaches the most life-critical systems; this stems from the frequent lack of adequate network segmentation in many hospitals. A human resources administration computer infected by a phishing email can serve as a bridge to eventually reach chemotherapy infusion pumps or ICU vital sign monitoring. Network segmentation—physically and logically separating life-critical devices from administrative systems and those with public or internet access—is a basic security measure that many hospitals have historically failed to implement because of cost, complexity, or lack of awareness.

Another limiting factor is the chronic shortage of qualified cybersecurity professionals in the sector. The global market for cybersecurity experts is extremely competitive, and public or nonprofit hospitals generally cannot offer the salaries and benefits of banks, technology consulting firms, or large software companies. Hospital IT teams are often overwhelmed in keeping legacy systems running, with no time or resources to devote to proactive security, threat analysis, or responding to sophisticated incidents. Partially outsourcing security functions to specialized firms can help, but it introduces its own trust risks and does not completely replace the need for in-house staff with deep knowledge of the specific clinical environment.

The consequences of a successful cyberattack on a hospital go far beyond direct economic damage or institutional reputation. First and foremost are the human lives that could be lost or harmed. In 2017, the WannaCry ransomware attack, which exploited a known window vulnerability, paralyzed dozens of National Health Service hospitals in the United Kingdom, canceling thousands of scheduled medical appointments, elective surgeries, and urgent diagnostic procedures. In 2020, a ransomware attack on a university hospital in Düsseldorf, Germany, forced the transfer of a critically ill emergency room patient to another facility more than 30 kilometers away; the patient died during the transfer, and German authorities formally investigated whether the delay caused by the cyberattack directly contributed to her death, which would have potentially constituted the first documented ransomware-related homicide. In 2021, a massive attack on the Irish healthcare system caused systems to be paralyzed for weeks, the cancellation of thousands of appointments, and the permanent loss of years of clinical data from patients who had not been properly backed up. Behind every aggregate statistic are individual patients who suffer real consequences from the disruption of their medical care.

Cyberattacks affect not only the availability of computer systems but also, more insidiously, the integrity of stored clinical data. A sophisticated attacker who gains privileged access could silently modify a patient's laboratory results, change a critical dosage in an electronic prescription, or subtly alter a diagnostic image, causing medical harm that is difficult to detect at the time and potentially lethal. If robust data integrity auditing mechanisms are not in place, it may never be known with certainty whether a clinical error was truly a human error or was maliciously induced by an attacker. This type of data integrity attack—more sophisticated and stealthier than visible ransomware—is also potentially more dangerous because it directly undermines the fundamental trust in clinical information upon which all medical decisions are based.

9.2.2 COMPREHENSIVE RESPONSE: TECHNICAL, ORGANIZATIONAL, AND CULTURAL

An effective response to this growing and evolving threat cannot be solely technical, nor can it be limited to simply purchasing more firewalls or antivirus software. It must be comprehensive, organizational, cultural, and, ultimately, clinically oriented. Hospital cybersecurity leaders must sit at the same coordination table as clinical and nursing leaders do, mutually understanding that decisions regarding IT security have direct and sometimes critical implications for the ability to care for patients, and vice versa. Clinicians and nursing staff must proactively understand that security measures are not unnecessary bureaucratic hurdles imposed by IT professionals out of touch with reality but rather essential protection for their patients' safety, analogous to hand hygiene or medication verification. Hospital executive leadership must allocate adequate financial and human resources to cybersecurity, explicitly recognizing that it is as important to

modern patient safety as traditional clinical safety or the prevention of nosocomial infections.

Staff awareness and ongoing training are key components of any defense strategy. A single careless click on a link in a sophisticated phishing email can initially compromise the entire corporate system. A weak or shared password can be the gateway for privilege escalation. A USB drive found in the parking lot and curiously plugged into a workstation could contain malware specifically designed for the healthcare sector. Training all hospital staff—from medical specialists and nursing staff to administrative personnel, lab technicians, and cleaning and maintenance staff—in basic cybersecurity hygiene practices is an investment with a guaranteed return on security. This training cannot be a one-time event at the start of employment; it must be ongoing, adapted to emerging threats identified in the environment, and periodically evaluated through phishing drills or incident response exercises.

Active collaboration among hospitals and between hospitals and regulatory and law enforcement authorities is also essential for effective collective defense. Cyber attackers share information about new vulnerabilities, successful attack tactics, and vulnerable targets in clandestine forums; legitimate defenders must do the same in authorized and trusted forums. Sharing technical indicators of compromise (IOCs), issuing timely alerts about newly detected phishing campaigns, and reporting incidents so that other hospitals can proactively protect themselves are practices that strengthen security across the entire healthcare sector. In many developed countries, healthcare-specific information-sharing and threat-analysis centers (ISACs) already exist that facilitate this necessary collaboration.

For medical software developers who create products deployed in these vulnerable hospital environments, understanding that the hospital is an implicit attack vector means designing with security as an absolute priority from the earliest stages of development, not as a late-stage add-on. In the design of medical software, one cannot assume that the hospital environment where the product will be deployed will be secure or benign; one must explicitly design with the assumption that the environment will be hostile, that there will be active attackers and that the software must withstand them. Medical data must be encrypted at all times, communications between components must be mutually authenticated, access must be thoroughly audited, and the software must be architecturally updatable because threats are constantly evolving and vulnerabilities are inevitably discovered after initial deployment.

For clinicians who use these systems daily, understanding this reality of persistent threats means accepting that cybersecurity is an inherent part of their modern professional responsibility. Reporting suspected security incidents, following established access policies, and actively participating in security training are not optional activities that can be ignored because of a lack of time. They are as much a part of their fiduciary responsibility to patients as they are to wash hands between patients, verify the patient's identity before administering medication, or maintain the confidentiality of information during conversations. Because a cybersecurity breach can have clinical consequences as serious as a traditional medical error, even if the pathways to harm are different.

The future of hospital cybersecurity inevitably hinges on the widespread adoption of “zero-trust” architectures, in which no device, no person, and no network connection is considered trustworthy by default without continuous verification. It involves comprehensive network segmentation, widespread encryption of data in transit and at rest, mandatory multifactor authentication for all privileged access, and continuous, automated monitoring of behavioral anomalies on the network. Effective public–private collaboration among hospitals, device manufacturers, software developers, and authorities is needed. It involves the development and adoption of specific, rigorous security standards for connected medical devices. Above all and ultimately, it involves widespread cultural recognition that modern cybersecurity is indistinguishable from traditional clinical safety and that protecting IT systems directly protects patients who trust us with their care.

9.3 SECURITY BY DESIGN, NOT BY PATCH

There is a way of building software—unfortunately still prevalent in many commercial development environments—that could be colloquially described as “functionality first, security second, if there's time and budget left.” In this superficial architectural approach, developers focus obsessively on quickly implementing the functional features the user has requested, on making the software “work” for normal use cases, and on meeting contractual delivery deadlines. Security is considered retrospectively at the end of the project as a superficial layer added on top of the

finished product: configuring a network firewall here, enabling some basic encryption there, and implementing rudimentary access controls in the user interface. This approach produces systems that are inherently insecure by design because security is not integrated into their fundamental structure but rather slapped on with digital duct tape as an afterthought. When a security vulnerability is inevitably discovered during operation, a patch is hastily applied; when another vulnerability appears, yet another patch is added. Over time, the system becomes a tower of accumulated technical patches, difficult to maintain operationally, impossible to fully audit, and surprisingly easy for a determined attacker to penetrate. This is patch-based security, and it is exactly what we must not do, under any circumstances—ly in healthcare environments where software security can literally be a matter of life and death for patients.

The mature and professional alternative is security by design: an architectural approach where security is considered from the very start of the project, as a nonnegotiable functional requirement, on the same level of importance as clinical functionality or system performance. In this approach, all architectural decisions are made with explicit consideration of their potential security implications. Threat models are systematically analyzed before the first line of production code is written. Security testing is continuously integrated into the development process, not as an isolated final milestone. Security is not something hastily added at the end but rather something that permeates the entire digital product, from its architectural foundations to its end-user interface, as an emergent property of sound design. This approach, while requiring a greater initial investment of time and expertise, produces systems that are inherently more secure, more resilient to attacks, and surprisingly easier and cheaper to maintain in the long term because it avoids the technical debt of accumulated patches.

9.3.1 FUNDAMENTAL PRINCIPLES OF SECURITY BY DESIGN

Early and systematic identification of potential threats: Before beginning to build any component, the team must rigorously ask itself, Who might have the motivation and interest to attack this specific system? What stored or processed data are most valuable on the black market or for competitive intelligence? What technical and human entry points exist in the proposed architecture? What specific types of harm could a successful attacker cause to patients, the institution, or public trust? These questions are not abstract academic exercises; they must be answered as concretely and realistically as possible, and the documented answers must directly guide subsequent technical design decisions. A well-developed threat model, following methodologies such as the STRIDE or PASTA, identifies the most relevant risks and allows security measures to be prioritized on the basis of their probability and potential impact.

Layered defense in depth: It is not enough to rely on a single layer of security; no matter how sophisticated it may be; one must professionally assume that any individual layer may eventually fail and have additional layers of protection that kick in when that happens. An attacker who manages to bypass the network perimeter firewall should immediately encounter strong multifactor authentication requirements. If authentication is compromised through social engineering, there should be granular access controls based on the principle of least privilege. If access controls fail because of a configuration error, there should be a layer of comprehensive auditing to detect and respond to unauthorized access. Each additional layer of defense exponentially increases the difficulty for the attacker and provides multiple opportunities for early detection and timely response.

The principle of least privilege applied systematically: every human user, every software process, and every system component should be assigned only the permissions strictly necessary to perform its legitimate function, and absolutely nothing more. A primary care physician needs access to the records of their own assigned patients but not to those of all patients in the hospital. An application for viewing lab results should not have permissions to modify them. A batch process that runs in the background to generate statistics should not have network access unless explicitly needed. Applying this principle consistently drastically reduces the total attack surface and limits the potential damage that a single compromised component can cause because its capabilities are restricted by design.

Security by default in configuration: the initial default configuration of any system should be as secure as possible, not the most open, permissive, or convenient for the administrator. If the system has configuration options that are inherently insecure or risky, they should be disabled by default until an experienced administrator explicitly enables them, with full, documented knowledge of the risks being assumed. This prevents systems that are hastily deployed in hospital

environments from being exposed because of carelessness or ignorance. It also greatly facilitates long-term security management because the secure configuration is factory default and does not require specialized intervention to enable it.

Separation of critical responsibilities: particularly sensitive functions should not rest with a single person or a single unsupervised technical component. The professional who authorizes privileged access should not be the same one who subsequently audits that access. The developer who writes the code should not be the only one who tests it from a security perspective. The systems administrator who manages the infrastructure should not be the same person who defines organizational security policies. This separation of duties hinders internal fraud and undetected abuses of power and provides independent cross-checks that increase the system's overall reliability.

Proactive minimization of the attack surface: Every implemented feature, every exposed interface, every open network port, and every service running continuously constitute a potential entry point for a creative attacker. For this reason, anything not strictly necessary for clinical operations must be removed or disabled. If an advanced feature is used by only a few specialized users, it should be disabled by default and require explicit activation. If a network service is not missioncritical, it should not run at all. If a remote administration interface is not routinely used, it should be closed and only opened temporarily when necessary. A minimalist, clean system is inherently easier to secure than one overloaded with features that no one uses but that leave doors open.

Secure and controlled error and exception handling: When a system experiences a failure, its error messages should not reveal technical information that could be useful to an attacker in refining their strategy. Error messages visible to the end user should be generic and noninformative regarding the internal structure, while full technical details should be directed solely to internal audit logs, which are protected and analyzable. A message that simply says "incorrect credentials" is secure because it does not distinguish whether the user exists or not; one that says "user does not exist" reveals valuable information about which identifiers are valid in the system. Errors must always be handled via secure code paths that do not leak sensitive information, and robust mechanisms must be designed for the secure recovery of a consistent state.

Secure and verifiable updatability: No software system is perfect, nor will it remain free of vulnerabilities forever; new weaknesses will always be discovered after initial deployment. For this reason, it is absolutely essential that the system be architecturally designed to be securely updatable (), without introducing new vulnerabilities into the update mechanism itself, without interrupting critical clinical service more than strictly necessary, and with cryptographic guarantees that the update is authentic and intact and has not been tampered with in transit by an attacker. Remote update mechanisms must be protected by robust digital signatures and must allow for gradual canary deployments that minimize the impact of potential issues not detected during testing.

Resilience to attacks and failures: Medical systems must be explicitly designed to continue operating—albeit in a degraded but safe manner—even when under active attack or after suffering partial compromises. This requires redundant architectures, isolation of critical components, and rapid, automated recovery capabilities. A system that goes completely down in the face of a distributed denial-of-service (DDoS) attack is not a secure system; it is a fragile system that has failed in its design. Security is not just about preventing successful attacks but also about withstanding them when they occur, degrading them gracefully, and recovering them quickly without loss of data or essential functionality.

Operational transparency and complete auditability: A truly secure system must leave detailed, immutable traces that allow one to retrospectively determine what occurred, who performed which action, and exactly when. These audit logs must be technically protected against subsequent tampering and must be analyzable proactively to detect anomalous behavior patterns before serious damage occurs. However, auditing is not just a forensic tool for use after an incident; it must also be used proactively in real time to identify suspicious behavior, unusual access, or patterns suggesting partial compromise of the system.

9.3.2 THE CULTURE OF SECURE DEVELOPMENT AND SHARED RESPONSIBILITY

For software developers working in the healthcare sector, consistently applying these ten principles in their daily work requires a profound and sustained shift in mindset. Security is not

something delegated to a specialized team of “security experts” who review the code at the end; it is the shared responsibility of all members of the development team. Every line of code must be written with active consideration of how it could be exploited by a resourceful and motivated attacker. Every architectural design decision must be explicitly evaluated from the perspective of its security implications. This does not mean being paranoid or paralyzed by risk analysis; it means being professionally rigorous. In other traditional engineering disciplines, safety is a natural and inseparable part of the work: a civil engineer does not design a public bridge without thoroughly considering its structural resistance to earthquakes, extreme winds, or unforeseen overloads. In medical software engineering, it should be exactly the same.

For clinicians and healthcare professionals involved in information system development or procurement projects, understanding these fundamental principles of security by design helps them evaluate and support the technical decisions of development teams, even when these involve apparent operational drawbacks. When a developer explains that a requested feature cannot be implemented quickly but insecurely because of security principles, they are not making technical excuses or being stubborn; they are professionally applying secure design principles that protect patients. When the system requires additional multifactor authentication to access particularly sensitive data, it is not an unnecessary bureaucratic hassle; it is an extra layer of defense that can save lives. Understanding these fundamentals improves interdisciplinary collaboration and reduces mutual frustration between clinicians and technical staff.

Security by design is neither an optional luxury nor a market differentiator; it is an absolute necessity in an environment where cyberattacks are becoming increasingly sophisticated and frequent and where the consequences of a security breach can be literally fatal for individual patients. Hospitals and healthcare systems cannot afford systems that are “secure” only until a determined attacker puts them to the test; they need systems that are architecturally designed to withstand attacks, to detect intrusions early, to respond effectively, and to recover quickly. This is achieved only if security is present, considered, and prioritized from the very first design sketch, in every line of code written, in every architectural decision made, and in every code review performed. It is more costly in terms of time and expertise at the beginning of the lifecycle but is infinitely less expensive than the catastrophic consequences of a security breach in production that affects real patients.

9.4 AUTHENTICATION, AUTHORIZATION, AND ACCESS AUDITING

“At the heart of healthcare IT security lie three eternal questions: Who are you? What can you do? And what have you done? The correct answers to these questions separate patient care from the chaos of intrusion.”

Let us imagine a completely ordinary scenario in any modern hospital during a typical work morning. An internal medicine physician accesses the electronic health record system to urgently check the laboratory test results of a patient who has just arrived at the clinic with acute symptoms. A ward nurse carefully records the administration of intravenous medication in the electronic care summary of a postsurgical patient. An admission clerk updates the contact information for a patient whose address has recently changed. A medical researcher requests formal access to anonymized data from a patient cohort for an epidemiological study approved by the ethics committee. A chronic care patient checks the results of their latest follow-up blood test from home via the hospital's portal. Behind each of these seemingly simple and routine actions lies three fundamental security questions that any robust healthcare information system must be able to answer accurately and reliably: Who exactly is the person or system attempting to access the data? What specific permissions do they have to perform the requested action on that specific data? Once they have accessed the data and performed operations, what permanent and immutable record remains of what they have done? These three universal questions correspond to the three inseparable pillars of access security in any sensitive information system: authentication, authorization, and auditing.

Authentication is the technical process of reliably verifying the identity of who or what is attempting to access the system. It is the concrete answer to the fundamental question, “Who are you truly?” In the healthcare setting, where identity confusion can have serious clinical consequences, this process must be particularly robust and multifactorial because the correct identification of users is critical to ensuring that sensitive information reaches precisely those who are legitimately authorized to access it and absolutely not those who are not.

Authentication methods have evolved significantly from simple, memorable passwords—

which have proven insufficient against multiple attack vectors—toward more secure and resilient multifactor authentication (MFA) schemes. Multifactor authentication combines something the legitimate user knows (traditionally a complex password) with something they physically possess (a registered mobile phone, an access smart card, a hardware token) or something they inherently are (a unique fingerprint, a facial recognition pattern, voice geometry). This combination of factors is rapidly becoming the de facto standard for accessing critical systems in hospital settings because the compromise of a single factor does not allow unauthorized access.

Authentication is not only necessary for human users who interact directly with the system. Automated systems that communicate with each other over networks—such as when a clinical laboratory sends results electronically to the central medical record system—also need to authenticate each other to establish trust. Both systems must be cryptographically certain that they are actually communicating with the entity they believe it to be and not with an imposter or an intercepted system. This is achieved through X.509 digital certificates issued by trusted certification authorities, shared and periodically rotated API secret keys, or other robust machine-to-machine authentication mechanisms. In a modern hospital environment with dozens or even hundreds of interconnected specialized systems, managing these digital identities in a consistent, scalable, and secure manner is an operational challenge that requires mature, specific infrastructure, such as centralized directory services (LDAP, Active Directory) or modern federated identity management (IAM) platforms.

9.4.1 GRANULAR AUTHORIZATION AND SPECIAL CONTEXTS

Authorization is the second essential pillar, which comes into play once the requester's identity has been successfully verified through authentication. It answers the critical question: "What specifically are you allowed to do in this system with these data?" Not all authenticated users have—or should they have—the same access and modification permissions. A treating physician can view and modify the medical records of patients under direct care but should not be able to change their own system credentials or access human resource data. A nurse can record vital signs and medication administration but should not be able to modify primary medical diagnoses. A patient can view their own personal data and clinical results, but not those of other patients or the physician's confidential notes. A researcher can access aggregated and anonymized data for their approved study but not individually identifiable information. Authorization rigorously defines these access rules on the basis of organizational policies, legal regulations, and professional ethics and enforces them technically in an automated and unavoidable manner.

The most common and traditional model for implementing authorization in information systems is role-based access control (RBAC).(6) In this classic model, individual users are assigned to predefined functional roles (doctor, nurse, administrative staff, patient, and researcher), and specific permissions for accessing resources are assigned collectively to these roles rather than to individuals. When an authenticated user attempts to perform an action on a resource, the system checks whether any of the roles assigned to that user explicitly have permission for that action in that context. This model is relatively simple to manage at scale and easy to understand conceptually, but it can be overly rigid and lacking in granularity in complex clinical environments where real-world situations define simple categorizations.

Owing to this limitation, more flexible and context-aware authorization models are emerging and maturing, such as attribute-based access control (ABAC), which considers multiple factors simultaneously when authorization decisions are being made: the user's nominal role, the operational context (time of day, physical location, device used), the type and sensitivity of the requested data, and whether or not there is an established therapeutic relationship with the patient whose data are being accessed.(7)

In the specific healthcare setting, authorization presents unique characteristics that make it particularly complex to design and implement correctly. One of these characteristics is authorization by temporary delegation: a physician may explicitly authorize a specific nurse to perform certain advanced tasks on their behalf during their absence, such as modifying medication doses within an established protocol—a delegation that must be revocable and auditable. Another critical feature is authorization for medical emergencies: in extreme clinical situations where delay can cost lives, it may be necessary to temporarily allow access that would not normally be authorized for that professional in that context while always leaving a clear, auditable record to ensure that there was no abuse. Another is granular authorization at the field level: it is not enough to state generically that a physician "can view records"; it must be technically specified that they can view

the records of their assigned patients but not those of patients from other departments and that within each record, they can view certain clinical data but not others that are particularly sensitive (for example, they can view general medical diagnoses but not detailed notes on psychiatry or mental health, which require specific additional authorization).

Auditing is the third fundamental pillar, and it answers the retrospective yet crucial question: “What exactly did you do with the system, and when?” Every access to sensitive identifiable data, every modification of a clinical record, every failed authentication attempt, and every privileged administrative operation must be permanently recorded in audit logs that allow for the subsequent reconstruction of events with forensic precision. These audit logs must include sufficient and detailed information, such as who accessed the data (authenticated identity), what specific data (resource identifier), from which network location and device (IP address, hostname), exactly when (precise timestamp with time zone), with what result (success, failure, denial), and what specific operation was attempted or performed. The logs must be technically protected against subsequent tampering so that not even system administrators with elevated privileges can modify or delete them without leaving an immutable trace of such an attempt. They must be efficiently analyzed, preferably with automated correlation and anomaly detection tools that identify suspicious patterns without continuous human intervention.

9.4.2 THE TRIAD IN ACTION: INTEGRATION AND CONTINUOUS MONITORING

Auditing is not merely a forensic tool for use after a security incident occurs; it also serves an important preventive deterrent function. Knowing that all access is permanently logged and reviewable deters potential nosy or malicious insiders from unjustifiably accessing data that are not concerned with them. In many audited hospitals, cases have sadly been detected among healthcare professionals who improperly access the medical records of famous patients, personal acquaintances, or even ex-partners, simply out of morbid curiosity—serious privacy violations that auditing allows them to be detected, documented, and subject to disciplinary action.⁽⁸⁾ Auditing also enables the technical investigation of confirmed security incidents, precisely identifying how an external attacker gained entry, which compromised the credentials they used, and what specific data they may have viewed, copied, or modified during their intrusion—information essential for regulatory reporting and mitigation.

Detailed audit logs in a hospital of this size generate massive volumes of data that defy manual analysis. A large hospital can easily generate millions of audit events per day, each with multiple metadata fields. Managing this massive amount of information requires specialized security information and event management (SIEM) systems—⁽⁹⁾—that centralize logs from multiple heterogeneous sources, normalize them into standard formats, and apply sophisticated temporal and contextual correlation rules to detect emerging threats. These SIEM systems are capable of identifying complex patterns that would be impossible for a human operator to detect manually: a specific user accessing an unusually high number of patient records in a short period of time, access from a geographic location unusual for that regular user, multiple consecutive failed authentication attempts followed by a suspicious success, or the correlation of seemingly innocuous events across different systems that together suggest a coordinated attack pattern.

The functional relationships among authentication, authorization, and auditing are intrinsically close and must be designed and implemented in an integrated manner from the architecture level, not as separate silos. Authentication provides the verified identity that authorization uses as input for its permission decision, and auditing comprehensively records both the results of authentication processes (successes and failures) and the authorization decisions applied, as well as the specific actions subsequently performed on the resources. A failure or weakness in any of these three links potentially compromises the integrity of the entire security system. For this reason, it is absolutely critical that medical software developers understand in depth how these three pillars are related to one another and implement them coherently and consistently in their applications, using frameworks and libraries that facilitate this proper integration.

For healthcare software developers, robustly implementing these three fundamental mechanisms requires selecting and properly configuring mature technologies and proven standards. For authentication, widely adopted standards such as OAuth2 for delegated authorization, OpenID Connect for token-based federated authentication, and SAML for integration with enterprise legacy systems exist, which enable secure integration with centralized corporate identity providers. For authorization, there are standardized policy languages such as

XACML (eXtensible Access Control Markup Language) or, more recently, JSON-based policies (JSON Web Tokens with authorization claims) that allow for the definition of complex, contextual, and automatically evaluable access rules. With respect to auditing, standardized formats such as Syslog or CEF (Common Event Format) facilitate integration with corporate SIEM systems.⁽¹⁰⁾ However, technology is only an enabler; what matters most is designing the right security policies from the outset on the basis of a careful and documented analysis of actual clinical requirements, applicable legal obligations, and patients' ethical expectations.

For clinicians and healthcare professionals who use these systems daily, understanding these security mechanisms helps them grasp why the system sometimes requires them to authenticate themselves with a second factor or why, under certain circumstances, they cannot access certain data even though they legitimately believe they should be able to. It also helps them be aware that all their access actions are permanently logged and that misguided curiosity or negligence in safeguarding their credentials can have serious disciplinary consequences, in addition to ethical violations. It provides them with the basic technical vocabulary needed to participate constructively in organizational discussions about access policies, contributing their invaluable expert knowledge about which accesses are truly legitimate and necessary in each specific clinical context.

A particularly delicate aspect of design in the healthcare domain is access management in acute medical emergencies, where security must not hinder the speed of clinical response. When a patient arrives at the emergency room in critical condition and is unconscious, the on-call physician needs immediate access to their prior medical history even if they are not the patient's regular physician or have an established therapeutic relationship. Information systems must incorporate technical "break the glass" mechanisms that allow for these exceptional urgent accesses but with compensatory controls to prevent systematic abuse. Typically, these emergency mechanisms require healthcare professionals to provide a documented clinical justification at the time, generate immediate automatic alerts to security and privacy supervisors, and leave a detailed record of the entire transaction that must be reviewed later by an audit committee. A delicate and necessary balance between the urgent clinical need to save lives and the fundamental protection of patient privacy is needed.

The correct integration of these security mechanisms with the patient's real identity is another critical technical challenge. When a patient accesses their own personal data through a hospital's web portal or mobile app, they need to authenticate themselves securely and verifiably, but it is also absolutely essential to ensure, both technically and organizationally, that they see only their own complete data—not that of another patient, which might appear due to a configuration error or software bug. This requires unequivocally and verifiably linking the patient's digital identity (their portal account) to their unique clinical identifier in the medical records database and ensuring through rigorous testing that this link is correct and persistent. Errors at this integration point can lead to serious and widespread privacy breaches, such as a patient mistakenly viewing another patient's complete medical record, with devastating legal and reputational consequences for the institution.

The future of authentication, authorization, and auditing in the healthcare sector points toward greater technical and contextual sophistication. Continuous, passive biometric authentication, combined with user behavior analysis (behavioral biometrics), promises more robust identity verification that is simultaneously less intrusive to the clinical workflow. Dynamic authorization based on real-time operational context and policies evaluated using artificial intelligence will enable more granular, adaptive, and context-appropriate access control.

AI-powered auditing will detect anomalous behavior patterns that today go completely unnoticed amid the noise of millions of events. However, these promising technological advances must not make us forget the fundamental point that the security of access to healthcare information is, ultimately and beyond technology, a matter of human trust. Patients trust that their most intimate data are protected by competent and ethical professionals. Healthcare professionals trust that they can access the information they legitimately need when they need it to provide better care. Society's trust that the healthcare system as a whole treats personal health information with respect to, care, and protection it deserves because of its unique nature.

Once we have established robust mechanisms to control who accesses what information and to leave an immutable record of it, how do we transform all that accumulated information ethically and securely into useful medical knowledge through statistical analysis and artificial intelligence, without betraying the trust of the patients who entrusted us with their data?

REFERENCES

- Introduction - Beyond the HIPAA Privacy Rule - NCBI Bookshelf [Internet]. Washington (DC): National Academies Press (US); 2009 [cited 2026 May 13]. Available from: <https://www.ncbi.nlm.nih.gov/books/NBK9576/>
- GDPR vs HIPAA: Key Differences & Compliance 2026 [Internet]. [cited May 13, 2026]. Available from: <https://www.atlassystems.com/blog/gdpr-vs-hipaa>
- Protection of Sensitive Data: What Companies Need to Know – Nunes & Asociados Study [Internet]. [cited May 13, 2026]. Available at: <https://estudionunes.com.ar/proteccion-de-datos-sensibles-lo-que-las-empresas-deben-saber/>
- Personal Data Protection Law: Key Points for Businesses [Internet]. [cited May 13, 2026]. Available at: https://www.innovaciondigital360.com/cyber-security/proteccion-de-datos-personales-en-argentina-claves-principales/#Ley_Nº_25326_de_Argentina
- Regulations on Personal Data: Ministry of Finance [Internet]. [Cited May 13, 2026]. Available at: [https://www.hacienda.gob.es/es-ES/El Ministerio/Paginas/DPD/Normativa_PD.aspx](https://www.hacienda.gob.es/es-ES/El%20Ministerio/Paginas/DPD/Normativa_PD.aspx)
- Coyne EJ, Weil TR. ABAC and RBAC: Scalable, flexible, and auditable access management. *IT Prof.* 2013;15(3):14-6.
- Hu VC, Ferraiolo D, Kuhn R, Schnitzer A, Sandlin K, Miller R, et al. Guide to Attribute-Based Access Control (ABAC) Definition and Considerations. January 2014 [cited May 13, 2026]; Available from: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-162.pdf>
- UCLA Hospitals Receives \$865K HIPAA Fine for Failing to Protect Celebrity Medical Records [Internet]. [cited May 13, 2026]. Available from: <https://www.hipaajournal.com/ucla-hospitals-receives-865k-hipaa-fine-failing-protect-celebrity-medical-records/>
- Security information and event management (SIEM) systems | Internal Revenue Service [Internet]. [cited May 13, 2026]. Available from: <https://www.irs.gov/privacy-disclosure/security-information-and-event-management-siem-systems>
- Ferraiolo D, Chandramouli R, Kuhn R, Hu V. Extensible access control markup language (XACML) and next generation access control (NGAC). *ABAC 2016 - Proc 2016 ACM Int Work Attrib Based Access Control colocated with CODASPY 2016* [Internet]. March 11, 2016 [cited May 13, 2026];13-24. Available at: [doi/10.1145/2875491.2875496?download=true](https://doi.org/10.1145/2875491.2875496?download=true)