AI for Resilient Health Systems: An Information Infrastructure Perspective*

Paolo Spagnoletti[1], Aysel Alizada[1]*, Francesco Nesta[2]

[1]*Department of Business and Management, LUISS Guido Carli University, Rome, Italy*
[2]*Department of Computer, Control, and Management Engineering Antonio Ruberti (DIAG), Sapienza University, Rome, Italy*

Abstract

National health systems increasingly rely on digital infrastructures to ensure continuity of care, support public health decisions, and coordinate clinical activities across organisational boundaries. As artificial intelligence becomes embedded in these infrastructures, resilience can no longer be understood solely as the outcome of human work. Rather, AI introduces new resilience capacities while simultaneously creating novel governance, interoperability, and regulatory challenges. This paper asks: How does AI reshape resilience in national health information infrastructures? Drawing on the information infrastructure literature, we conceptualise resilience as a multi-level capacity of the health system, generated in part through the evolution of its digital infrastructures, at three complementary levels: (1) epidemiological resilience through population-level analytics, (2) operational resilience through interoperability and continuity of information flows, and (3) clinical resilience through AI-assisted decision support for healthcare professionals during clinical care.

We investigate these dynamics through an interpretive case study of the evolution of Italy's national electronic health record (Fascicolo Sanitario Elettronico, FSE 2.0) and Health Data Ecosystem (Ecosistema Dati Sanitari, EDS) between 2022 and 2026. The study follows two distinct regulatory trajectories of artificial intelligence: the first, proposed in the 2022 EDS scheme, was not permitted within that legal and architectural context and was subsequently removed from the 2024 decree following the opinion of the Italian Data Protection Authority; the second is represented by MIA, a clinician-focused AI assistant designed to support Italian general practitioners and other healthcare professionals, which aligns with the subsequent legislative framework that provides specific safeguards and implementing measures. Rather than treating these episodes as isolated implementation successes or failures, we analyse how they reveal the institutional, technical, and organisational conditions under which AI successfully supports, or is actively constrained from supporting, the resilience of national health systems.

The paper contributes to information infrastructure research by extending existing accounts of resilience beyond human practices and by demonstrating that resilience increasingly depends on the governance of AI capabilities embedded in evolving digital infrastructures. In practice, the study offers design implications for implementing the European Health Data Space and for other national initiatives seeking to integrate AI into large-scale health information infrastructures.

Keywords: Information Infrastructures, Artificial Intelligence, Resilience, Electronic Health Records (FSE 2.0), Health Data Ecosystems (EDS)

1. Introduction

Safety has traditionally been defined as a condition in which nothing goes wrong. Since it is impossible to guarantee that nothing will ever go wrong, safety is often described as a condition in which the number of adverse events remains acceptably low, although the term 'acceptably' is itself ambiguous. The term 'resilience' was introduced to represent a new perspective on safety and was rapidly recognised as a valuable complement to established approaches in both

The 12th International Conference on Socio-Technical Perspectives in IS (STPIS'26), September 16–17, 2026, L'Aquila, Italy.

*Corresponding author.
aalizada@luiss.it (A. Alizada); 0009-0007-8065-9751

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-10
© 2026 Copyright for this paper by its author(s). Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR-WS.org/Vol-4325/paper36.pdf

industrial and patient safety. Scholars and practitioners in industry and academia acknowledged that resilience offered a structured framework for addressing challenges related to complexity, interconnectedness, 'system of systems', and ultra-high reliability environments. Over time, it has become recognised that resilience extends beyond only managing threats and disturbances or addressing situations where failures may occur. Currently, resilience is more broadly defined as a system's intrinsic capacity to adjust its functioning before, during, or after changes and disturbances, thereby sustaining essential operations under both expected and unexpected conditions. This perspective highlights not only the capacity to maintain operations, rather than solely react to and recover from disruptions, but also the continuous capacity to adapt to changing conditions and, where possible, sustain and improve system performance [53].

This conceptualisation of resilience raises an important question about how systems operate and adapt over time. In information systems research, this issue has traditionally been addressed through the concept of maintenance, which is understood as an ongoing activity involving the correction, adaptation, and improvement of operating systems. Building on Swanson's differentiation among corrective, adaptive, and perfective maintenance [98], Edwards [51] expanded the scope from software maintenance to systems maintenance, defining it as modifications to system elements necessary to eliminate faults, adapt to changes within the system or its components, or enhance performance. Additionally, Edwards distinguishes between reactive maintenance, which addresses identified changes, and proactive maintenance, which involves monitoring changes and assessing their effects across system elements [51]. This distinction matters for resilience, as a system's ability to adjust its functioning depends not only on its capacity to recover after disruptions but also on its ability to detect and accommodate changes before they compromise essential operations.

Introducing IS in healthcare is complex because healthcare providers (with diverse professional roles, training, and experience), patients (with diverse conditions, personal characteristics, and medical trajectories), and medical treatments (with diverse procedures and approaches) have different needs. A system that works well in one setting may fail in another or require different reorganisations [61]. National health information infrastructures play a key role in making healthcare more connected and efficient. The national health information infrastructure is defined as a set of technologies, standards, applications, systems, values, and laws that support all facets of individual health, health care, and public health [48]. It includes an information network that allows data to flow in three areas: (1) personal health, to help people manage their own wellness and health care choices; (2) health care providers, to give them access to complete and accurate patient information at any time and to tools that help with clinical decisions; and (3) public health, to handle and monitor public health issues and health education programs [48]. Electronic health records (EHRs) are a central component of this infrastructure, as they collect and make available patients' health data over time and across different care settings [49, 50]. Together with other technical components and governance arrangements, EHRs help develop a national health information infrastructure that improves the effectiveness and efficiency of healthcare delivery [49]. National health information infrastructures should therefore be understood as evolving sociotechnical arrangements rather than finalised technological systems.

Artificial intelligence adds additional complexity, as AI components are integrated into existing infrastructures rather than functioning independently. These components rely on data, standards, interfaces, organisational processes, professional practices, and regulatory frameworks. As a result, implementing AI can change both the AI component and the broader infrastructure.

The Italian electronic health record (Fascicolo Sanitario Elettronico, FSE 2.0) and the new national data platform (Ecosistema Dati Sanitari, EDS) were both designed to include automatic data analysis features. Two distinct regulatory encounters with AI characterise the evolution of Italy's national health information infrastructure. First case: the 2022 draft scheme for the EDS decree proposed automatically processing FSE data for the entire population and explicitly relied on artificial intelligence systems [3, 35]. The national data protection authority (Garante) issued a negative opinion on 22 August 2022 [35], noting that the scheme failed to regulate the conditions for lawful processing and did not specify the security measures required to mitigate risks of re-identification, profiling, and automated decisions affecting individuals. The authority's favourable opinion of 26 September 2024 pertained to a substantially revised EDS, in which it declared the automatic processing of FSE 2.0 data inadmissible and removed all references to AI systems from the immediate deployment mandate [3, 30, 4]. The use of AI within the EDS was not prohibited but postponed to a subsequent regulatory reform that would specify purposes, legal bases, controllership roles, data subjects' rights, and adequate technical and organisational measures, in accordance with the GDPR, the AI Act, and European guidance

on the automated processing of health data. Second case: in January 2026, MIA ("Medicina e Intelligenza Artificiale"), a clinician-focused AI assistant developed by AGENAS, the National Agency for Regional Health Services, launched a twelve-month national pilot involving 1,500 general practitioners. The platform provides optional, guideline-based clinical support, with one clear restriction: during the pilot phase, it does not access patients' personal data [6, 33, 97]. Its design follows the framework established by Articles 7 to 10 of Law 132/2025. According to Article 7, artificial intelligence in healthcare may be used only to support prevention, diagnosis, care, and therapeutic decision-making; the physician remains responsible for the final decision, and the patient must be informed of its use. Articles 8 and 9 regulate the processing of health data for research, experimentation, and secondary purposes. Furthermore, Article 12-bis, introduced into Decree-Law 179/2012 by Article 10, establishes the AGENAS platform, with data types, operations, and technical and organisational measures to be determined following consultation with the Ministry of Health, the Garante, and the National Cybersecurity Agency [39, 97]. This platform is being deployed in three phases: analysis and design; testing of the primary care role by designated physicians; and long-term platform management and maintenance [5].

Beyond these two specific initiatives lies a broader, equally critical dimension of resilience: operational continuity. National health information infrastructures depend not only on population-level analytics and AI-assisted clinical decision-making, but also on the continuity and interoperability of information flows across regional healthcare organisations. In Italy, where healthcare is organised at the regional level, the FSE relies on a heterogeneous installed base that varies in digital maturity, organisational practices, and adoption patterns. Resilience therefore depends not only on whether AI can support population-level analytics or clinical decision-making, but also on whether the wider infrastructure can sustain and adapt the information flows on which healthcare operations depend.

Against this background, this research-in-progress paper asks: how does AI reshape resilience in national health information infrastructures? To address this question, the study combines an interpretive, longitudinal single-case study of the evolution of Italy's FSE 2.0/EDS from 2022 to 2026 [40, 41, 42] with the problematisation strategy proposed by Alvesson and Sandberg [47]. This question guides the analysis of resilience in three connected ways: epidemiological resilience through population-level analytics; operational resilience through interoperability and continuity of information flows across regions; and clinical resilience through AI-assisted decision support for healthcare professionals during clinical care. The non-admitted analytics layer and the conditional admission of MIA offer two real-world examples for studying epidemiological and clinical resilience. Operational resilience is studied by examining how well healthcare systems share information and collaborate across regions.

This research is timely because Regulation (EU) 2025/327 on the European Health Data Space (EHDS) requires all Member States to enable the secondary use of electronic health data, with key provisions applying from March 2029 [7, 8]. The EHDS also arrives alongside the NIS2 Directive, the AI Act, and the Cyber Resilience Act, so the admission of AI into health infrastructures is increasingly shaped by a convergent regulatory architecture rather than by data-protection law alone [78, 79, 80]. The Italian case therefore offers an opportunity to examine the organisational and governance challenges that other national health information infrastructures are likely to face as AI components become progressively integrated into health data ecosystems. The study aims to contribute both theoretically and practically. Theoretically, it examines how AI reshapes resilience across epidemiological, operational, and clinical levels within national health information infrastructures. In practice, it seeks to derive design implications for implementing AI-enabled health information infrastructures within the emerging European regulatory framework.

The paper is structured as follows. Section 2 provides the theoretical background, examining information infrastructures and resilience, the role of the installed base, standardisation, generativity in sustaining resilience, and the implications of delegating aspects of infrastructure maintenance to artificial intelligence. Section 3 outlines the research design and methodology. Section 4 presents a case study of Italy's FSE 2.0 and EDS, detailing infrastructure functionality, the integration of AI components, regulatory convergence influencing compliance, operational challenges highlighted by a significant cyber incident, the progression from design to non-admission and conditional enablement, a comparative analysis of the MIA platform, and the conditional pathway for AI-enabled secondary use under the EHDS. Section 5 discusses the study's theoretical and practical contributions, and Section 6 concludes.

.

2. Theoretical background

2.1. Information infrastructure and resilience

Previous studies have examined electronic health record systems as information infrastructures, showing how they evolve through installed bases, standardisation efforts, local adaptations, and distributed forms of governance [1, 9, 10, 2]. These studies have helped explain how national health infrastructures remain resilient despite heterogeneity and fragmentation. Within this literature, information infrastructure has three main dimensions**:** an installed base of users and practices [11], standards that enable interoperability [12], and generativity, which allows users to adapt the system locally [13]. The first two dimensions help an infrastructure last and stay stable, while the third supports innovation and new uses [14].

Hanseth and Lyytinen proposed a design theory to explain the complex evolution of Information Infrastructures (IIs). These are *shared, open, heterogeneous*, and *evolving* socio-technical systems built on Information Technology (IT) capabilities. This definition highlights both the structural and emergent properties that make IIs different from their individual parts (see Table 1). There are four main emergent properties: *Shared*, meaning an II is used by many different communities in ways the original designers did not expect. *Open*, which means there are no fixed boundaries, so new components, users, and designers can join at any time. Because of this, designs must assume that form, content, and access will always stay open. *Heterogeneous* — as IIs become more diverse over time, not just in technical parts but also in the types of people and groups involved, such as user communities, operators, standardisation and governance bodies, and design teams. *Evolving* — since IIs are open, they keep changing and never reach a final state. They are not built from scratch and do not end; instead, they grow through ongoing experimentation and by involving new communities, changing gradually rather than all at once [60].

Distributed control is typically the most effective approach to managing changes in information infrastructures (IIs), since a single authority cannot direct their evolution solely from the top down. Instead, it depends on the existing system, requiring each new component to integrate with established infrastructures, capabilities, learning processes, and shared conceptual frameworks [60].

Property	*Application*	*Platform*	*Information infrastructure*
1) Emergent			
Shared	Yes — locally, through specified functions	Yes — across the user communities involved and a set of IT capabilities	Yes — universally, across multiple IT capabilities [1, 67]
Open	No — closed by user group and functionality	Partially — depends on design choices and managerial policies	Yes — universally, allowing unlimited connections to new user communities and IT capabilities [68, 69, 70]
Heterogeneous	Partly, mainly through the social groups involved	Partly — mainly via social groups, but also technical connections	Yes — increasingly heterogeneous both technically and socially [59, 66, 71, 72]
Evolving	Yes, but bounded by time horizon and user community; linear growth; evolution bounded and context-free	Yes, but limited by architectural choices and functional closure; mostly linear growth; path-dependent evolution	Yes — unlimited by time or user community [1, 70, 73]; both linear and nonlinear growth [66]; path-dependent evolution [1, 59, 67]
2) Structural			
Organising principle	Direct composition of IT capabilities within one homogeneous platform	Direct composition of a set of horizontal IT capabilities within a set of homogeneous platforms	Recursive composition of IT capabilities, platforms and infrastructures over time [1, 59]
Control	Centralised	Centralised	Distributed and dynamically negotiated [68]; can rely only on basic organising principles (standards) and installed-base inertia [1, 59]

Table 1: Applications, platforms and information infrastructures. [1]

Hanseth and Lyytinen's design theory explains the structural properties [60]. Technically, an II is made up of other infrastructures, platforms, applications, and IT capabilities in a repeating pattern; socially, it also repeats, serving as both the result and the basis of design action. Control is shared and happens in stages, through negotiation and agreements that decide which groups manage which parts, what is connected, and who can access it [60].

This paper draws on resilience engineering, in particular the work of Hollnagel and his colleagues, who applied the concept to health care. They argued that adverse events often result from a combination of conditions, not just failures of single parts or human error. They also argued that failures are simply the opposite side of successes, so there is no need to look for special failure mechanisms. To build resilience in health care, we need to study how health systems actually work, rather than focusing only on how they fail [53]. From this perspective, resilience is based on four key abilities: responding to events, monitoring ongoing situations, learning from past experiences, and anticipating potential developments [99].

Health system resilience is the ability to prepare for and respond to crises, maintain essential functions during emergencies, and adapt based on lessons learned. Building resilience depends on the situation and is a continuous process. It involves assessing system strengths and weaknesses in advance, investing in vulnerable areas before a crisis, reinforcing the system during emergencies, and reviewing performance afterwards. For instance, how quickly a system recovers from a crisis is an important sign of success [52]. Recovery represents only one of several meanings that the term resilience has acquired. Woods identifies four concepts for resilience in the literature: resilience as rebound, defined as the return to previous functioning after a disruption; as robustness — the ability to absorb a known range of disturbances; as graceful extensibility — the capacity to stretch when unexpected events challenge a system's boundaries, which contrasts with brittleness; and as sustained adaptability — the ability of layered networks to continue adapting to future surprises as conditions evolve [100]. Woods contends that the first two concepts offer limited value. Rebound focuses on reactive phases, and recovery depends on pre-existing capacities before a surprise occurs. Robustness applies only to well-modelled disturbances, and expanding a system's ability to handle certain perturbations can increase its vulnerability to others. Two additional points are directly relevant to this analysis. First, organisations may inadvertently undermine their own sources of resilience by overlooking how individuals compensate for adaptive shortfalls, a risk that emerges when such responsibilities are delegated to AI. Second, sustained adaptability relies on governance that functions across layered networks over extended time scales, aligning with the layered and distributed nature of information infrastructures. Accordingly, this study understands resilience in the senses of graceful extensibility and sustained adaptability, rather than as rebound or robustness [100]. In national health information infrastructures, AI components built on standardised data encounter an installed base characterised by regional variation that is difficult to model. The central question is therefore not only how quickly the system recovers, but whether it maintains the capacity to stretch when variation exceeds its components' design parameters.

From an information-infrastructure perspective, much of this capacity depends on the health system's digital systems. Drawing on information infrastructure research, this study conceptualises resilience as a multi-level capacity of the health system, shaped in part by how its digital systems evolve [52, 53]. It distinguishes three levels: (1) epidemiological resilience, through population-level analytics; (2) operational resilience, through interoperability and continuity of information flows across regions; and (3) clinical resilience, through AI-assisted decision support for healthcare professionals during clinical care. These levels focus on what the health system can do, not only on the features of its technology. Together, these accounts complement each other: the four abilities and the concepts of graceful extensibility and sustained adaptability define resilience as a capacity, while the three levels indicate where within the health system this capacity is generated and observed. Section 3 examines the research questions at these levels.

2.2. Resilience through the installed base

Monteiro and Hanseth describe II as an open, shared, evolving, standardised, and heterogeneous installed base [14]. A key idea in this literature is that information infrastructures are never designed from scratch: they develop through the evolution of an installed base [2]. The installed

[1] Adapted from Hanseth and Lyytinen [60].

base, or “what is already there”, covers current work practices, tools, the division of labour, and legal and professional rules; it both helps and limits how infrastructures can change [2, 15]. Infrastructure does not grow de novo; it wrestles with the “inertia of the installed base” and inherits its strengths and limitations [1]. Infrastructures change gradually rather than all at once [16], which helps explain why parts of a national infrastructure can develop at different rates across regions. Because of this, the literature recommends shifting from a top-down control approach and instead focusing on cultivating the installed base [2]. A *cultivation approach* recognises the installed base and aims to make changes incrementally. *Growth* happens naturally in a changing environment through repeated changes, as infrastructure adapts to, changes, or includes parts of that environment [59]. A cultivation strategy has the following features: a) *process-orientation* means working step by step with technology and established practices, making gradual changes to the infrastructure over time; b) *user mobilisation* refers to how control is shared in a cultivation strategy; c) learning is about the choices made during the cultivation process. Cultivation is a learning-based strategy in which designers decide which parts work well and which do not [61]. The installed base includes the politics, standards, ontologies, timing, and ways people interact within the infrastructure [17]. Infrastructure stability relies on the decentralised actions of its base, which reveal how social and technological components connect over time and across places [16, 1, 18].

From a resilience perspective, the installed base serves a dual role: its established arrangements provide continuity, while its interdependencies and path dependencies can constrain adaptation. Inertia should not be equated with resilience. Instead, resilience refers to the installed base’s capacity to accommodate change while maintaining essential infrastructural functions. Heterogeneity, local adaptations, and decentralised arrangements may facilitate incremental adaptation, while increasing interdependencies can create lock-in effects that hinder future change. Cultivation may therefore support resilience by enabling incremental changes to the installed base, preserving operational continuity, and reducing the risk that accumulated dependencies become rigid forms of lock-in.

Large infrastructure systems show that managing them involves more than just centralised, top-down control. Unlike 'systems,' which are usually part of an organisation, infrastructure is a shared resource for a wider community. As a result, they cannot be designed or managed like isolated information systems. This requirement limits how new elements can be designed, which means the existing infrastructure, or installed base, strongly affects how the infrastructure will develop in the future [74, 75]. Hanseth and Braa describe the shift from a system to an emergent infrastructure as a four-step process. First, the application takes on an infrastructural role simply because of its scale. This includes the solution's size, the variety of units and functions it supports, and the number of users, managers, and developers involved. Second, the idea of a single universal system often fails. When organisations validate and implement the system, they usually customise it to fit local needs. This process turns the original application into different versions for each unit, resulting in many overlapping and connected applications that together form a complex and heterogeneous infrastructure. Third, this web connects to the main corporate infrastructure, and, through this connection, links to the Internet. Finally, over time, the installation becomes increasingly connected to nearby installations, so the parts can only develop together [15].

Unlike other design theories and methods that “design from scratch”, Hanseth and Lyytinen’s design theory focuses on the installed base: II development involves building a self-reinforcing installed base by using existing ones and avoiding the risks that come with them [60].

2.3. Resilience through standardisation

Experts agree that as standards and standardisation become more complex, old models no longer fit today’s challenges. Because the world is changing rapidly, standards must be more flexible to adapt [19]. Hanseth discusses two kinds of flexibility — use and change flexibility — and argues that standards need both [14]. A standard's total flexibility is the sum of these two. *Use flexibility* means being able to apply a standard in several different environments, or for several different purposes. The "need to know" principle, which focuses on gathering only essential information rather than "nice to know" information, shows how use flexibility can help to get the most information from the least data. *Change flexibility* comes from the classic principle of modularisation: instead of creating one complex standard for everything, it is better to create several small, simple standards and set up clear interfaces, or gateways, between them [19].

This stable situation is only part of what researchers have found. Monteiro and Hanseth showed in their study that information infrastructures need standards to exist. For communication to

happen, partners must use a common standard, which acts as a "language" or, more technically, a protocol. Standardised protocols allow computers from different vendors to communicate. This is a key part of a general-purpose, open information infrastructure [14]. David and Greenstein identify three types of standards: reference, minimum quality, and compatibility standards. II standards are compatibility standards that ensure a component can fit into a larger system by following interface rules. The compatibility standards for II create a complex network. Sometimes, standards are replaced, in whole or in part, by new ones, leading to a "genealogy of standards." The heterogeneity of II standards, how one standard includes or is connected to several others, is an important part of II [76].

Standards support infrastructure resilience by enabling interoperability and mitigating the negative effects of heterogeneity through compatibility expectations enforced by the installed base [19, 20]. Standards connect heterogeneous systems and help define the installed base's power, scope, and quality [21]. Once a standard diffuses, modifying it becomes difficult because diffusion creates lock-ins and self-reinforcing effects called technological momentum [22]. Like the installed base, standards thus contribute to the stable and efficient operation of an information infrastructure.

2.4. Resilience through generativity

Zittrain describes generativity as a technology ecosystem's ability to create new output, structure, or behaviour without input from its original creator. In other words, generativity reflects a technology's capacity for leverage across a range of tasks, adaptability to different tasks, ease of mastery, and accessibility [25]. Generativity provides an infrastructure that can adapt to new environments, incorporate elements from its surroundings [21], and operate in different places [23]. Infrastructures are never truly finished and can be used in ways not yet imagined [24]. Unlike highways or sewer systems, infrastructure shows generativity by enabling new outcomes without additional input from its designers. This is described as the "capacity to produce unprompted change driven by large, varied, and uncoordinated audiences" [25]. In health infrastructures, generativity comes from the system's data practices and services. It is created by the distributed installed base, including regional authorities, clinicians, local IT staff, and vendors, through practices like workarounds, local repairs, and interpretive data work [26, 23, 17]. These same practices also support resilience, as practitioners notice missing, uneven, or incorrect data and make adjustments.

Recent research finds that this generative capacity comes from architecture. Yoo, Henfridsson, and Lyytinen explain that digital innovation uses a layered modular architecture. Generativity occurs because the layers are loosely connected, allowing innovations to start in any layer and spread to others. Unlike a purely modular product, in which components fit into a single design hierarchy, a layered modular architecture allows components to belong to multiple hierarchies [62]. This idea draws on Simon's design theory, which argues that breaking a system into loosely connected parts with defined interfaces makes it less complex and more flexible [77].

II's resilience is the product of three different forms of innovation [61]. While innovation *of* infrastructures corresponds to conceptualising and implementing a new infrastructure, including reconceptualising and re-engineering existing infrastructures, innovation in infrastructures denotes replacing or modifying existing components of an infrastructure without changing its constituting architecture. Finally, innovation *on* infrastructures means extending existing infrastructures by adding new modules on top of (or in addition to) what exists. Infrastructure resilience then refers to adding new features or services to existing infrastructure, a process known as generativity. The three types of infrastructure are interconnected, and effective infrastructure strategies need to consider all three [61]. Hospital evidence confirms the architectural condition: Bygstad and Øvrelid show that process innovation succeeded when lightweight solutions at the clinical edge aligned with, yet remained architecturally distinct from, the heavyweight core through explicit governance and architecture-alignment mechanisms [64]. Generativity does not happen on its own; instead, platform owners regularly adjust the boundary resources others use to build [63], and ecosystems must remain both stable and capable of evolving [65].

2.5. Infrastructure maintenance and the AI delegation of adaptive work

Integrating AI-based components introduces distinct disruptions to this socio-technical equilibrium. Classic socio-technical design frameworks — most notably Mumford's ETHICS approach — assume that stabilising an information system requires ongoing, participative human intervention to balance technical capacity with social and

organisational needs [85, 86]. However, when an AI layer introduces autonomic self-healing features, this balance shifts.

From the infrastructure perspective, resilience relies on people recognising where the existing system is weak or uneven and finding ways to adjust. An AI component works differently: the missing data and small errors that humans used to manage carefully may instead now be absorbed and spread by its automated results, silently and at a national scale. Infrastructure theory was built for systems in which people performed this work; the introduction of AI, therefore, compels a re-evaluation of core infrastructure concepts when the human-centric assumption no longer holds.

Three assumptions remain unclear, each linked to one of the problematisation questions above, and each creating a researchable tension. First, maintenance: working infrastructure usually goes unnoticed until it breaks down [1]. Infrastructure studies describe its repair as ongoing human work [16, 26, 17]. An AI layer, however, can provide continuous and silent correction, fulfilling autonomic computing's goal of self-healing systems that automatically detect, diagnose, and repair their own problems, which is one of its four self-management properties [58] — *maintenance without maintainers* — so the moment at which the system recognises its own gaps may disappear. Second, the locus of judgment — that is, where and by whom decision-making is exercised: cultivation — design seen as stepping into ongoing processes (Section 2.2) — assumes flexible decision-makers who know when a rule should be adjusted [2]; automated parts replace flexibility with fixed limits, moving judgment from the edge of practice to the design phase. Third, trust: in established systems, trust in components builds up over time as they are used.

These tensions, rather than a lack of research, drive the research question. The first tension concerns repair-and-maintenance research, in which maintenance is ongoing, local, and mostly invisible human work that creates order rather than merely supporting it [54, 55]. When AI is added, it does not just automate this work, it changes how people understand it. The second and third tensions resemble findings from studies of AI in the workplace, which show that using AI changes, rather than replaces, expert work and shifts judgment and uncertainty between people and machines [56, 57]. In the Italian case, the three tensions appear at the three resilience levels introduced in Section 1. The maintenance tension applies at the operational level: automated controls and transformations in information flows currently make standardisation visible rather than invisible. The judgment tension appears at the epidemiological level, in the fixed-limit decision-making envisaged in the not-admitted analytics layer and the trust tension appears at the clinical level.

This delegation of human adaptive practices to automated layers redefines the core infrastructure. Bednar and Welch, writing in the same ETHICS tradition as Mumford, observe that introducing cognitive computing and advanced automation requires a shift from passive technical alignment to active, continuous co-evolution of social and technical elements, ensuring that systemic resilience is not compromised by opaque algorithmic controls [92].

AI-based components are assumed to include both back-end analytical systems that process aggregated data from national health infrastructures to produce population-level insights and front-end clinical decision-support tools deployed at the point of care. These components represent distinct modes of AI integration within health information infrastructures. The back-end systems generate automated outputs for population-level health planning, quality assessment, and care evaluation, while the front-end tools support clinical decision-making for individual patients. In both contexts, the analysis examines how AI interacts with the underlying information infrastructure, considering variations in data completeness, interoperability, and governance across regions.

3. Research design and method

Building on the theoretical tensions identified in the previous section and following the problematisation approach of Alvesson and Sandberg [47], this study investigates the central and related research questions:

RQ: How does AI reshape resilience in national health information infrastructures?

Rather than assuming AI adds new capabilities to existing infrastructures, the study examines how introducing AI interacts with the ongoing maintenance and adaptive work that sustain infrastructural functioning. In particular, it investigates three related questions:

RQ1: How does introducing AI reshape infrastructure maintenance?

RQ2: What happens to the resilience of an information infrastructure when part of the aaptive work traditionally carried out by humans is delegated to or automated by AI components?

RQ3: What new theoretical tensions emerge from this shift?

Figure 1 provides a visual summary of the interpretive, longitudinal single-case process study [40, 41, 42] of the evolution of the FSE 2.0/EDS between 2022 and 2026. The study examines AI-enabled resilience across three conceptual levels — epidemiological, operational, and clinical — that correspond analytically to the macro, meso, and micro levels, respectively. **Epidemiological resilience** at the macro level, which involves population-level analytics for public health and highlights the socio-technical tension between data-driven public health utility and data privacy protection, is examined through the non-admission of analytics capabilities for the secondary use of health data**. Operational resilience** at the meso level, which focuses on maintaining information flow and interoperability across facilities and highlights the tension between centrally mandated standards and locally established working practices, is examined through selected regional implementations and how AI may contribute to the continuity and interoperability of regional health information infrastructures. **Clinical resilience** at the micro level, which uses AI tools to support healthcare professionals during clinical care and highlights the tension between guideline-based automation and professional discretion, is examined through the conditional introduction of MIA as an AI assistant supporting primary care physicians.

Because the AI components under investigation are not yet fully embedded as operational components of the health information infrastructure, the study cannot study their organisational implications through direct observation. To explore how AI-enabled resilience is understood across different contexts, the planned interviews will combine a discussion of documentary evidence with three short scenarios corresponding to the macro, meso, and micro levels. These scenarios are not intended to predict future behaviour but to elicit participants' interpretations of how adaptive work is currently performed and how it might change when AI components become embedded within health information infrastructures.

Data analysis will proceed iteratively by combining the completed documentary analysis with the planned interview material. First, the documented sequence (design → non-admission → substitution → conditional admission) provides the initial process narrative and will inform the interview protocol. Second, interview transcripts will be analysed interpretively to identify how participants understand AI's role in shaping resilience across the macro, meso, and micro levels of the health information infrastructure. We will then interpret emerging themes through the theoretical lens of information infrastructure, with particular attention to concepts such as the installed base, standards, and generativity. This iterative evaluation of system changes and user feedback treats infrastructure evolution as a double-helix relationship between design architecture and organisational reflection-upon-use [94].

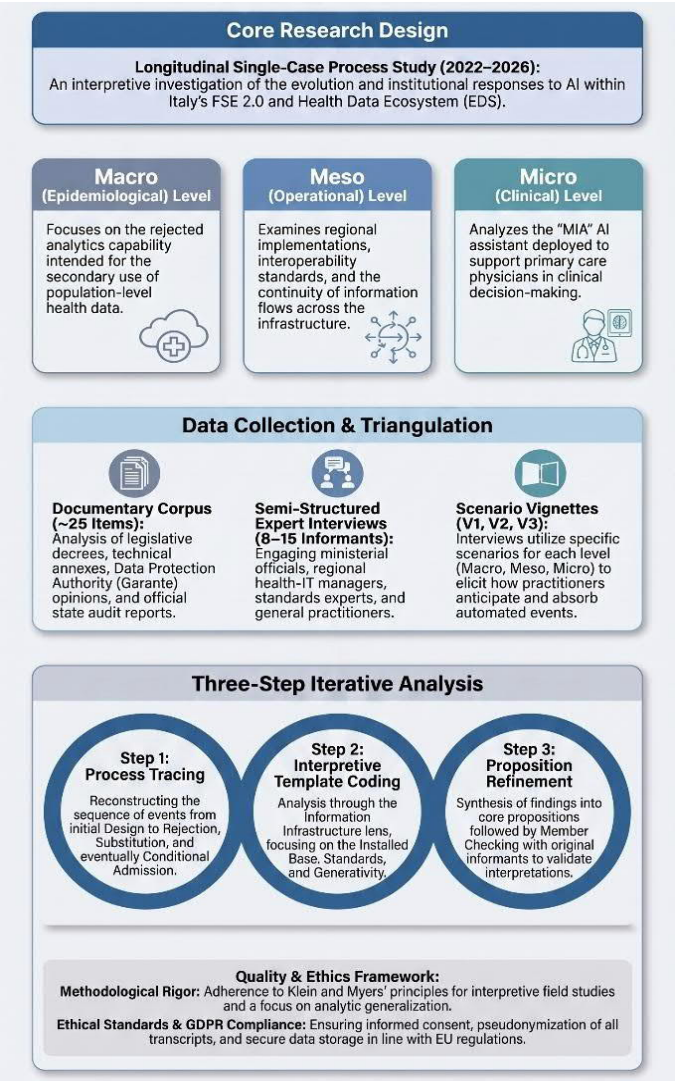


Figure 1: Research methodology overview

Finally, we will compare and refine insights from the interviews and documentary sources through iterative interpretation, including participant feedback on preliminary findings where appropriate.

Throughout the analysis, we will follow Klein and Myers' principles for interpretive field research [84] and seek analytic rather than statistical generalisation [40, 42].

This qualitative approach allows us to unpack the nuanced institutional power relations and shifting patterns of human-machine trust that characterise large-scale

digital infrastructures, aligning with contemporary frameworks for critical, interpretive information systems research [90]. The study will involve expert interviews concerning policy and organisational practices and will not involve patient data or access to clinical records. All participants will provide informed consent, interview transcripts will be pseudonymised, and research data will be stored and managed in accordance with applicable GDPR requirements.

4. Case study: Italy's FSE 2.0 and the EDS

4.1. How the infrastructure functions

Within the Italian National Recovery and Resilience Plan (PNRR), M6C1 and M6C2 are the two primary components of Mission 6, which focuses on Health. While M6C2 is most closely linked to innovation, research, and digitalisation, [43] notes that M6C1 also supports "e-Health" through its focus on proximity networks, local structures (such as Community Houses), and telemedicine. Mission 6 Component 2 (M6C2) specifically funds the technological evolution of the Electronic Health Record (FSE) and strengthens national digital health infrastructure [43]. The FSE is defined as the *"cornerstone" (pietra angolare)* for delivering digital health services and enhancing national clinical data. It is a primary focus of Mission 6 of the National Recovery and Resilience Plan (PNRR), which allocates billions to modernise and digitalise the National Health Service (SSN) [43].

The FSE 2.0 reform (Ministerial Decree of 7 September 2023) changes the national electronic health record system from document sharing to structured, standardised data sharing [27, 28, 29]. Each time a patient receives care, the provider sends the clinical documents and data to the regional record. There is one FSE for each region and autonomous province. Standardised information models make the data interoperable and easy to interpret [4, 29].

The EDS (Ecosistema Dati Sanitari) was created as part of Mission 6 of the national recovery plan (PNRR) by the Ministerial Decree of 31 December 2024. It is a national platform that uses data already collected in the FSE 2.0 for care, prevention, scientific research, and governance, including health planning, quality-of-care assessment, and healthcare evaluation [30, 4, 31]. The EDS is a single infrastructure that keeps the same health data in three different forms, each managed by specific groups, for different purposes and under different rules. These are not just steps in a process but three separate parts, each with its own data flow. Data are created when a healthcare facility produces a clinical document during a medical service and then sends it to the Electronic Health Record (FSE). A regional Gateway converts it to the standard FHIR format, extracts key data, and verifies compliance with standards before the data moves on [4, 29, 32]. These operations run through software procedures and algorithmic rules, which alone do not qualify the Gateway as an artificial intelligence (AI) system; a functional description of a component's actions is insufficient to associate an algorithm with AI. The Gateway operates within a broader layer of enabling services, including master patient index, terminology services, identity and access management, and application programming interface (API) gateways. The transition from document to data involves a structured sequence of validation, indexing, semantic mapping, and persistence, with each transformation requiring versioning and traceability to its source. Without such semantic governance, data may be technically accessible but not reliable, comparable, or reusable. Transparency at this layer therefore includes validation rules, schema versions, terminology catalogues, transformations, error codes, provenance, and logs. From there, the data take two paths: one goes to the cleartext regional units, and the other, after pseudonymisation, goes to the single national pseudonymised unit [30, 3]. A Broker manages requests between these units but does not store any data itself [30, 3, 4]. **Figure 2** shows the three-level architecture of the EDS:

- **Level 1 — Clear Data:** contains clear-text identifiable personal health data (e.g., names and clinical reports). It is stored in twenty-two physically separate archives: twenty-one regional units *(UA-R)* and one for seafaring personnel. A patient's data resides in the regional unit where they are assisted, and a Broker facilitates communication between these archives without storing any data itself. Access is restricted to healthcare professionals for treatment purposes and requires explicit consent.
- **Level 2 — Pseudonymised Data:** direct identifiers are replaced by unique tokens, remaining classified as personal data since identity can be reconstructed using a decryption key. Unlike Level 1, these data are stored in a single archive *(UD-P)* with logical partitions for each Region. The Ministry of Health, Agenas, and Regions have access for planning and quality assessment. Given the public interest, explicit consent is not required; however, cross-referencing data to re-identify individuals is strictly prohibited.

- **Level 3 — Anonymised Data:** this data is no longer personal, and identity cannot be reconstructed. The *UD-A* unit is not a permanent archive; instead, it provides "anonymisation on demand," creating temporary "views" for specific research queries. Agenas manages access and evaluates projects from public and private researchers. No consent is required.

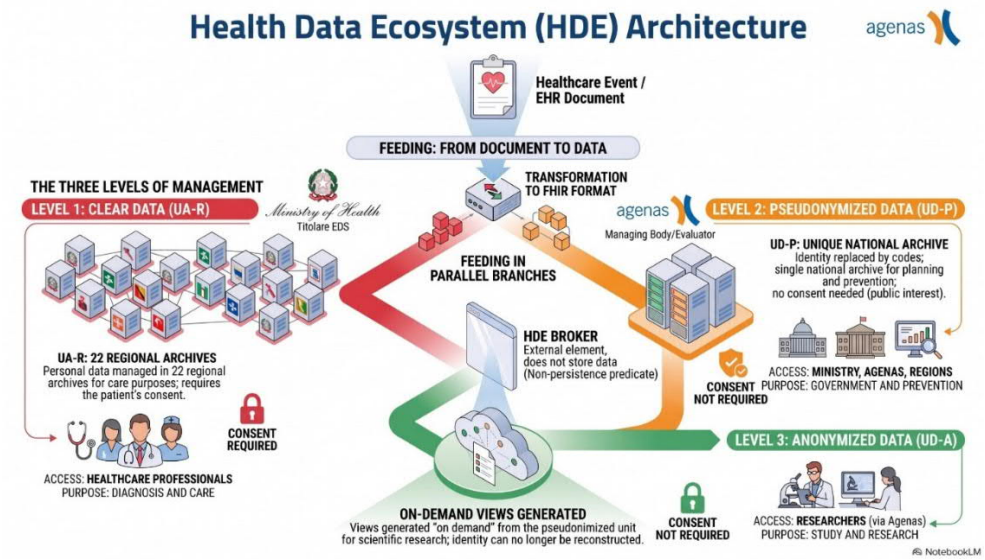


Figure 2: The three-level architecture of the EDS. Source: own elaboration.

Figure 3 illustrates the data flow of the Italian National Telemedicine Platform (PNT), comprising 21 Regional Telemedicine Infrastructures (IRT) and one National Telemedicine Infrastructure (INT), in which clinical information generated at the regional level is progressively de-identified before reaching the national layer. The flow begins when the producing system sends the telemedicine document to the regional technological solution for validation, which forwards it in clear form to the regional Electronic Health Record (FSE) and, in parallel, extracts its structured data for transmission to the Health Data Ecosystem (EDS). The EDS pseudonymises these data in accordance with the EDS decree, and Agenas then accesses its services without data persistence to process the pseudonymised records through the INT, whose internal modules compute anonymous indicators published on monitoring dashboards. The INT also exposes enabling services, such as the business glossary and the telemedicine solutions manager, that do not process personal data. In this architecture, identifiable documents and data remain confined to the regional layer, while the national layer operated by Agenas works exclusively on pseudonymised inputs transformed into non-reidentifiable indicators used for telemedicine tariff setting, health governance and programming, quality-of-care monitoring, and Health Technology Assessment, thereby reconciling national oversight with the GDPR principle of data minimisation.

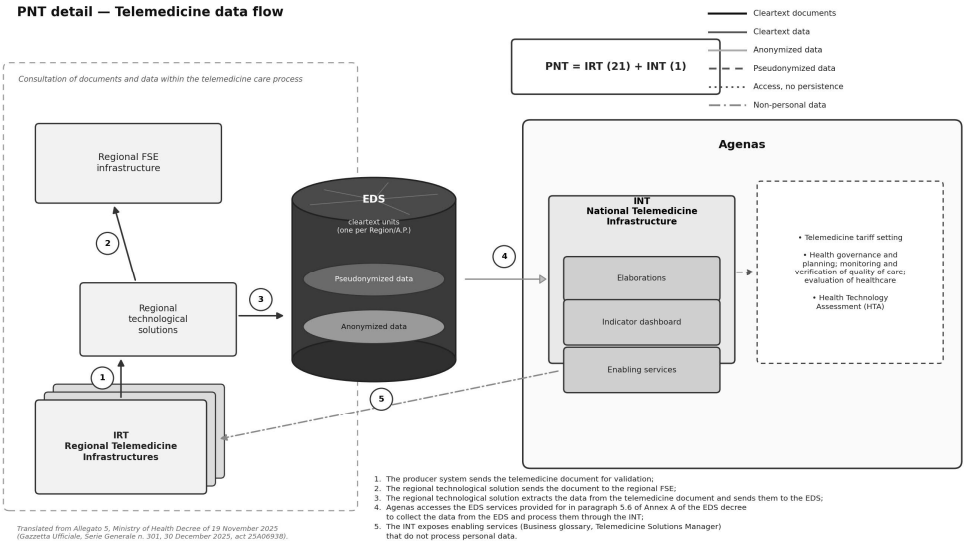


Figure 3: Telemedicine data flow through the FSE 2.0/EDS architecture [83].Source: Ministry of Health[2]

In Italy, the FSE relies on a heterogeneous installed base that varies in digital maturity and local organisational practices. This reflects what Bednar and Welch [93] characterise as inescapable contextual dependencies, where large-scale systemic mandates must contend with unique, locally historically shaped operational realities.

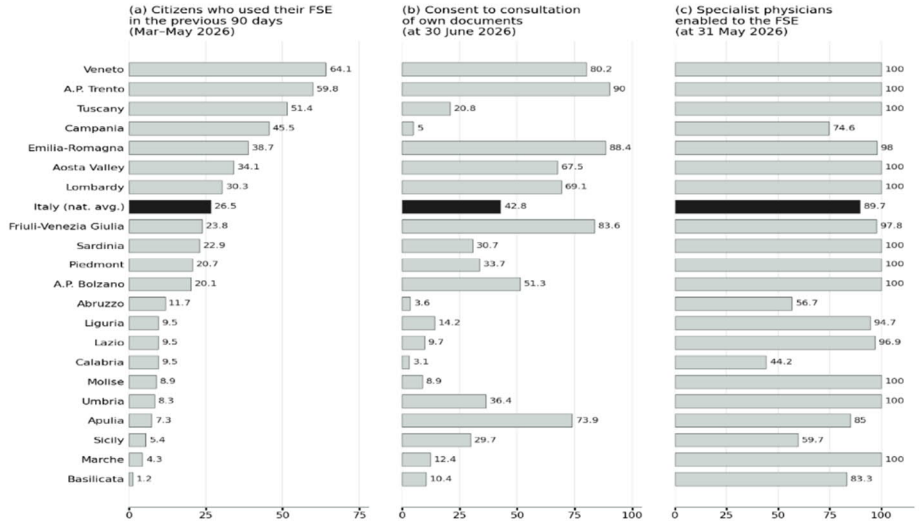


Figure 4: Regional distribution of three FSE 2.0 indicators from the national monitoring portal. Source: Ministry of Health[3]

Recent monitoring data provide a clearer view, showing that the installed base operates at two speeds, with regional differences that go beyond a simple North-South split. **Figure 4** shows the regional distribution of three FSE 2.0 indicators from the official monitoring portal [82]: a) the percentage of citizens who used their FSE in the previous 90 days (March to May 2026) ranges from 64.1% in Veneto to 1.2% in Basilicata, with a national average of 26.5%; b) citizens who gave consent for their documents to be consulted (as of 30 June 2026) varies from 90.0% in the Autonomous Province of Trento to 3.1% in Calabria, with a national average of 42.8%; (c) specialist physicians in public health authorities enabled to use the FSE (as of 31 May 2026) shows that supply-side enablement is almost complete – 11 regions have reached 100%, and the national average is 89.7%, but Calabria is again the lowest at 44.2%.

Two points are important to understand. First, the three measures do not rank regions the same way: Tuscany has high use (51.4%) but low recorded consent (20.8%), Campania has high use (45.5%) with the second-lowest consent (5.0%), while Apulia is the opposite (73.9% consent, 7.3% use). This shows that differences in the installed base are complex, not just a simple North–South divide. Second, the 2021 pattern continues in a new way: the infrastructure is mostly ready, but people's use is much slower, so any AI system using FSE data would still face very uneven data feeding and use [82].

4.2. AI components in and around the infrastructure: types, functions, status

The integration of AI within the infrastructure can be classified into three distinct categories rather than a single class of components **(Figure 5):** two documented AI initiatives, one algorithmic component supporting interoperability, and potential future scenarios for the secondary use of health data. (i) The EDS analytics layer, the first documented AI initiative, was introduced by a 2022 decree to enable automated processing of FSE data and AI-based insights for governance, but was subsequently debated and removed (Section 4.4) [3, 35, 4]. (ii) Gateway functions serve as an algorithmic interoperability component rather than an AI system. The Gateway ensures the correct and homogeneous formation of documents and data feeding the EDS by applying formal and semantic controls to documents and corresponding data produced by health and social-health facilities for the FSE, and by extracting that information into structured data in established standard formats [4, 29]. Although the state audit institution reported experimenting with AI algorithms as part of the Gateway's milestone releases, it did not make public any technical details about these algorithms or their operational functions [32].

[2] Ministero della Salute, Gazzetta Ufficiale della Repubblica Italiana, Serie Generale No. 301, 30 December 2025, https://www.gazzettaufficiale.it/eli/gu/2025/12/30/301/sg (accessed 2 October 2026)

[3] Ministero della Salute, "Fascicolo Sanitario Elettronico 2.0 – Utilizzo del Fascicolo Sanitario Elettronico," https://monitopen.fse.salute.gov.it/usage (accessed 2 October 2026).

Based on the available technical documentation, these functions cannot be classified as AI applications without additional technical evidence: the Gateway automates controls and transformations, making part of the infrastructural work of standardisation visible, but does not demonstrate the adaptive or inferential capabilities that characterise AI. (iii) MIA, the AGENAS point-of-care platform and the second documented AI initiative, is a decision-support tool that supports general practitioners by answering clinical questions and suggesting tests, prescriptions, and preventive measures based on certified clinical guidelines; during its trial phase, it does not access patients' personal data (Section 4.5) [5, 33, 34]. (iv) Future secondary-use AI, a prospective scenario rather than a deployed component: these are models that will be trained on pseudonymised and anonymised EDS data as part of the European Health Data Space, starting in 2029 (Section 4.6) [7, 8].
Among these, types (i) and (iii) are the two documented AI initiatives and the main focus of this study. Type (ii) is included as infrastructural context and type (iv) as a prospective horizon.

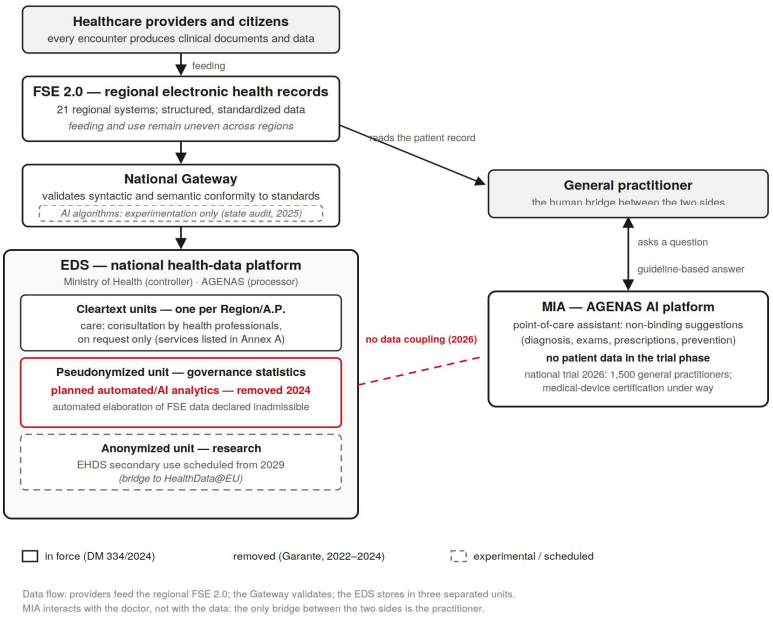


Figure 5: The location of AI components. Solid elements are in force; the red mark is the analytics layer removed in 2024; dashed elements are experimental or scheduled. Source: own elaboration.

Viewing through a multilevel lens, each component occupies a different level of the infrastructure: the analytics layer (i) and future secondary-use models (iv) sit at the epidemiological level, the Gateway functions (ii) at the operational level, and MIA (iii) at the clinical level. **Figure 5** shows where these components sit within the architecture, while **Figure 6** places the three documented trajectories on a single timeline, highlighting the study's main objects.

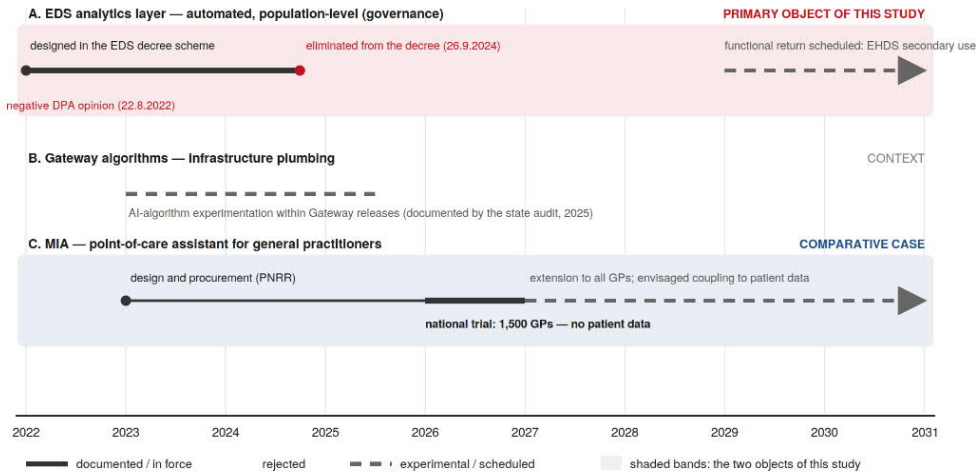


Figure 6: Regulatory trajectories of AI and the algorithmic role of the Gateway in the Italian healthcare information infrastructure (2022-2031). Source: own elaboration.

Operational-level resilience is a real concern. On 3 May 2023, a ransomware attack by the Monti group hit the local health authority ASL 1 Abruzzo: booking systems stopped working, staff returned to pen and paper, and about 500 gigabytes of health data — including records of HIV-positive and cancer patients — were stolen and later published. The national cybersecurity agency called this one of the most serious recent cybercrimes [45]. This event shows what happens when multi-organisational digital workflows collapse, reinforcing McEvoy and Kowalski's [95] socio-technical warning that cyber risks are fundamentally socio-technical. Vulnerabilities materialise not merely from technical flaws, but from the systemic degradation of daily working practices over time, which ultimately compromises an infrastructure's capacity to contain operational shocks. From this perspective, using AI to adapt at this level can both protect and create new risks, depending on whether it strengthens or further weakens the working practices that keep the system secure. The incident further illustrates the socio-technical reality documented by Bednar, Welch and Sadok, whose study of 471 employees across 259 organisations found that formal, centralised security mandates are often disconnected from employees' daily work practices: most could not even locate their organisation's security policy, and existing policies were typically outdated and written without reference to actual work context [96]. When infrastructure layers fail to account for local human routines, technical controls are bypassed in practice rather than followed, leading to systemic vulnerability and an inability to withstand operational shocks. Since then, the region has set up one of central Italy's first IoT-cybersecurity labs in Pescara [46].

The Abruzzo ransomware disruption underscores that infrastructure recovery is fundamentally an organisational and human effort, not a purely technical one [45]. This directly aligns with Holmström et al.'s socio-technical framework for operationalising cyber resilience [89], which argues that system durability during security shocks emerges from the tight coupling of social practices with automated infrastructure controls. It analyses this through four interdependent capabilities — anticipate, withstand, recover, and adapt — each assessed against the principle of joint optimisation: whether a capability's technical and social components are proportionate to the situation or imbalanced. The pen-and-paper fallback demonstrates that withstand capability is maintained almost exclusively through social mechanisms once the technical system fails, mirroring the imbalance observed in Holmström et al.'s example of staff unable to operate redundant servers [89]. In contrast, the Pescara laboratory represents an adapt-capability response implemented after the incident, rather than a solution enacted during the event [46]. This situation reflects not merely technical failure, but resilience sustained through social compensation, which is precisely the imbalance that their framework is designed to expose.

4.3. Regulatory convergence and infrastructural compliance

Alongside the GDPR, which guided the Garante's actions, four new frameworks are being introduced between 2023 and 2029: the **NIS2** Directive on cybersecurity risk management for essential and important entities, which Italy has adopted through Legislative Decree 138/2024 [78]; the **AI Act** sets requirements for risk management, data governance, event logging, human oversight, robustness, and post-market monitoring for high-risk AI systems. Clinical decision support falls into this category because it is linked to medical-device law [79]; the **Cyber Resilience Act** addresses the security of products with digital elements throughout their life cycle, including vulnerability management and software bills of materials. Certified medical devices are not included here, as they are already covered by medical-device regulations [80, 81]; the **EHDS** has been in effect since March 2025 and will be introduced in steps starting in 2027, with the first important phase for main uses and most secondary uses planned for March 2029 [7].

Although these instruments differ in their focus and target audiences, they share key features: risk-based governance, clear responsibility, supply-chain security, traceability, and proof of compliance. **Figure 7** illustrates how these elements align on a timeline from 2023 to 2029 which consists of seven distinct stages: 1) the Garante's 2023 Decalogue established the conditions for providing national health services using artificial intelligence systems [91]; 2) the 2024 EDS decree and the subsequent favourable opinion were issued following the exclusion of AI-based processing [3, 30]; 3) in 2025, the EHDS entered into force, and Law 132/2025 established the legal foundation for MIA [7, 97]; 4) the MIA trial initiated in 2026, aligning with the implementation of the Digital Omnibus [6, 88]; 5) in 2027, the EHDS became generally applicable, alongside the Annex III high-risk regime and the Cyber Resilience Act [7, 79, 80]; 6) from 2 August 2028, regulations concerning AI systems incorporated in regulated products take effect [79]. Their applicability to MIA depends on the platform's final classification under the

MDR/IVDR and the necessity of third-party conformity assessment [81]; 7) the initial operational phase of the EHDS begins on 26 March 2029 [7].
The matrix categorises each relationship between a regulatory regime and a subject matter into three degrees of proximity: **direct and primary impact**, where the subject matter is the explicit and central focus of the regime; **significant impact**, where the regime imposes requirements or produces material consequences for the subject matter without it being the principal focus; and **indirect or cross-cutting impact**, where effects arise from coordination with other regimes, infrastructural interdependence, or systemic consequences, without specific obligations directed at that subject matter.

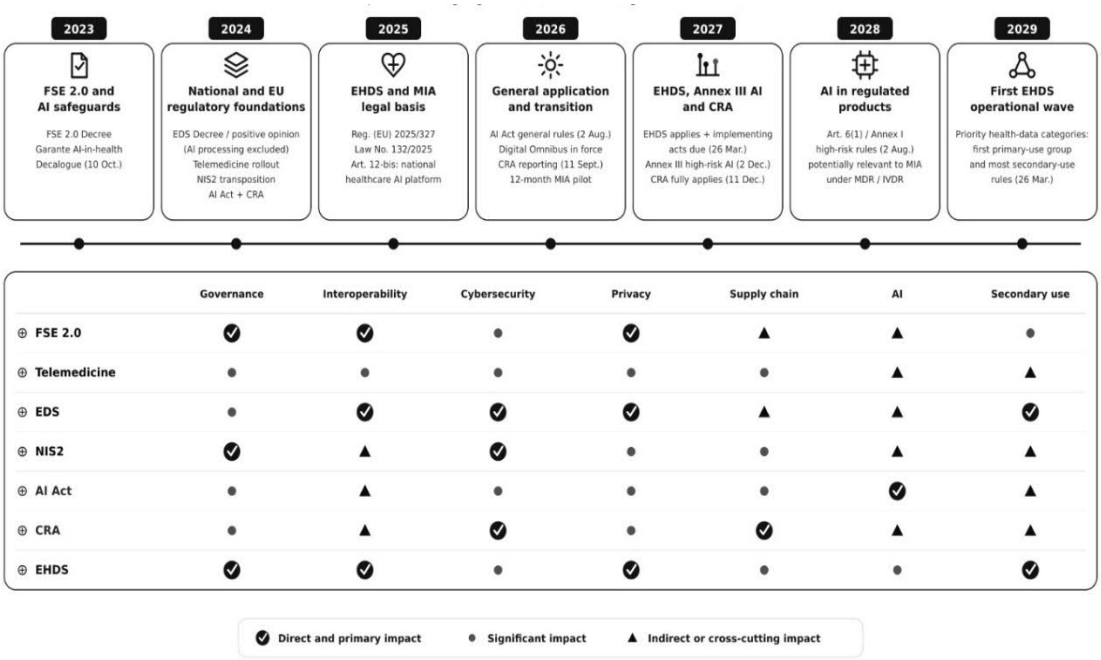


Figure 7: Converging national and EU regulatory and infrastructural milestones bearing on national and regional health-data infrastructures, 2023–2029. Source: own elaboration.

European regulations now require compliance to be built into systems, not just documented separately. For example, the GDPR calls for data protection by design (Art. 25), the AI Act includes its requirements in the design of high-risk systems [79], and the Cyber Resilience Act does the same for product security [80]. For this condition, we use the term *"infrastructural compliance,"* in which each legal requirement is clearly linked to a technical control, a specific responsibility, and verifiable evidence throughout the system's life cycle. In this context, conformity tools such as impact assessments, certificates, release checks, and audit logs help regulators and platform owners determine what to include in the infrastructure [63, 65]. The Italian case demonstrates this process, as discussed in Section 2.4: the Garante's conditions on the EDS, MIA's certification as a medical device under the MDR [81], and required practitioner training.
Two more mechanisms put compliance into practice. First, obligations pass to external suppliers through contracts, such as the data-processing agreements required by Art. 28 GDPR and the supply-chain security measures required by NIS2 [78]. Second, obligations are verified before release: high-risk AI must be classified, undergo a fundamental rights impact assessment alongside the DPIA, and maintain logs and technical documentation under the AI Act [79], while non-device software must meet the vulnerability-handling and software bill-of-materials requirements of the CRA [80].

Regime	Core obligations (selected)	Bearing on resilience levels in this study
GDPR and national FSE 2.0/EDS decrees	Lawful bases, purpose limitation, data protection by design, DPIA, Garante opinions.	Epidemiological: conditions for secondary use; grounded the 2022 non-admission and 2024 substitution.

NIS2	Cyber risk management, incident reporting, supply chain security, and management accountability for essential entities.	Operational: continuity of regional information flows; the 2023 Abruzzo attack as the stress case.
AI Act	For high-risk systems: risk management, data governance, logging, human oversight, robustness, post-market monitoring; FRIA coordinated with DPIA.	Clinical: governs the admissibility of decision support, such as MIA, and any future coupling with patient data.
Cyber Resilience Act	Secure development, vulnerability handling and disclosure, SBOM, support periods (reporting from Sept. 2026; general application Dec. 2027); certified medical devices excluded.	Operational: life-cycle security of non-device software-Gateway components, APIs, integration services.
MDR/IVDR	Conformity assessment, clinical evaluation, cybersecurity within the general safety and performance requirements, post-market surveillance.	Clinical: the regime under which MIA is certified as a medical device.
EHDS	Primary/secondary-use regimes, Health Data Access Bodies, data permits, secure processing environments (first phase March 2029).	Epidemiological: governs the conditional pathway for AI-enabled secondary use (Section 4.6).

Table 2: Converging regulatory regimes and their bearing on the three levels of resilience. Source: own elaboration

Table 2 summarises this section. The right column shows our analysis, indicating which of the three resilience levels in this study — *epidemiological, operational, or clinical* — each regime affects in the Italian case: data-protection law applies at the epidemiological level, NIS2 and the CRA address the operational level, the AI Act and the MDR cover the clinical level, and the EHDS plans for the epidemiological return of secondary use.

4.4. From design to non-admission to conditional enablement

This component developed in five distinct steps, collectively indicating not a prohibition of AI in healthcare, but a transition from an inadequate legal and regulatory framework to subsequent conditional regulatory enablement.

(1) **Design:** as the National Data Protection Authority (Garante) subsequently observed, the 2022 plan would have established the nation's largest health database, aggregating identifiable patient documents on a large scale without pseudonymisation mechanisms [3]. The draft scheme also proposed automatically processing FSE data for the entire population and explicitly relied on artificial intelligence systems [3, 35].

(2) **Non-admission**: on 22 August 2022, the Garante issued a negative opinion on an EDS scheme that lacked sufficient clarity on the data and services involved, the roles of actors, the conditions for lawfulness, algorithmic risks, and corresponding safeguards. The objections included generating automated clinical insights without an explicit request from a treating professional, delegating essential architectural rules to subsequent decrees, resulting in an 'empty box,' and inadequately defined data quality protocols [35]. The authority also noted that the scheme failed to specify adequate security measures to mitigate risks of re-identification, profiling, and automated decisions affecting individuals [35]. This episode demonstrates that the admissibility of AI in healthcare depends on the presence of specific regulatory provisions and safeguards, and that the absence of such measures, rather than the technology itself, was the basis for non-admission.

(3) **Systematisation:** in September 2023, the Garante published a Decalogue outlining the conditions for the delivery of national health services using artificial intelligence systems. These conditions include the legal bases for data processing, accountability and privacy by design and by default, clarification of the roles of involved actors, principles of knowability, non-exclusivity in algorithmic decision-making and algorithmic non-discrimination, requirements for data protection impact assessments, data quality, integrity and confidentiality, fairness and transparency, human supervision, and considerations related to personal dignity and identity that extend beyond data protection law [91].

(4) **Substitution**: the positive opinion issued on 26 September 2024 was about a redesigned EDS that no longer included the AI component, and a decision was made only after the automatic processing of FSE 2.0 data was declared unacceptable, services were limited to those explicitly requested and fully specified in an annex, and all mentions of AI systems were removed from the immediate deployment mandate. Approval also depended on fully implementing the

updated FSE 2.0 rules following the 2023 opinion [3, 36, 30]. Consequently, advanced population-level analytics were legally decoupled from the initial nationwide EDS deployment [37]. The use of AI within the EDS was not prohibited, but postponed to a subsequent regulatory reform specifying purposes, legal bases, controllership roles, data subjects' rights, and adequate technical and organisational measures, in accordance with the GDPR, the AI Act, and relevant European guidance on the automated processing of health data [3, 30, 4].
(5) **Legislative enablement**: Law 132/2025 authorises the use of artificial intelligence in healthcare, subject to defined purposes and limitations, human supervision, reliability standards, and specific implementing measures (Articles 8 to 10 and Article 12-bis of Decree-Law 179/2012, which establishes the MIA platform) [39, 97]. More specifically, Article 7 permits AI in healthcare only to support prevention, diagnosis, care, and therapeutic decision-making; the physician remains responsible for the final decision, and patients must be informed when AI is used. Articles 8 and 9 regulate the processing of health data for research, experimentation, and secondary purposes [39]. Article 12-bis establishes the AGENAS platform, with the categories of data, operations, and technical and organisational measures to be defined following consultation with the Ministry of Health, the Garante, and the National Cybersecurity Agency [39, 97].

4.5. A comparative trajectory: the MIA platform

Article 10 of Law 132/2025 introduced Article 12-bis into Decree-Law 179/2012, regulating the use of artificial intelligence in the healthcare sector and, specifically for territorial care, establishing an artificial intelligence platform managed by AGENAS. Under this framework, in January 2026, the AGENAS platform MIA ("Medicina e Intelligenza Artificiale"), developed under the PNRR [34], began a 12-month national trial involving 1,500 general practitioners selected by the Regions, due to conclude on 31 December 2026 [6].The platform is being deployed in three phases: analysis and design; testing of the primary care role by designated physicians; and long-term platform management and maintenance [5].
In its current configuration, the platform allows physicians to submit questions and receive responses grounded in certified knowledge, including guidelines and diagnostic-therapeutic pathways, as well as non-binding recommendations on examinations, prescriptions, and prevention strategies.
MIA does not replace the physician's decision-making authority and, during the current phase, does not process patients' personal data while undergoing certification as a medical device and mandatory practitioner training [5, 33]. These features are characteristic of the trial phase rather than a fully implemented system, while the complete configuration of the platform—including the services offered, categories of authorised users, types of data processed, operations performed, and security measures adopted—is delegated to the acts specified in paragraphs 1, 3, and 5 of Article 12-bis. MIA's definitive characteristics will be formalised in these acts, which are anticipated at the conclusion of the trial period [39, 97]. The comparative case is therefore examined while its admission terms are being determined, highlighting that AI integration into a national health infrastructure is established incrementally through successive regulatory acts, rather than being fixed at deployment.

4.6. The conditional pathway: secondary use under the EHDS

Regulation (EU) 2025/327 on the European Health Data Space requires each member state to allow the secondary use of electronic health data, with most rules taking effect in March 2029 [7, 8]. Italy is getting ready: in May 2026, the Ministry of Health established the national steering committee (*Cabina di regia*) for EHDS implementation [38], and Law 132/2025 facilitated the reuse of de-identified health data for research [39]. The current EDS architecture therefore does not include the advanced population-level analytics removed from the initial deployment mandate. Instead, the EHDS creates a future regulatory pathway for the secondary use of health data. That pathway is conditional rather than automatic: any future deployment of AI would still require a new design, a permitted purpose, a legal basis, a data permit or other prescribed procedure, a secure processing environment, and compliance with the GDPR, the EHDS Regulation, the AI Act, and national law [7, 8, 39]. The system's initial experience with AI can help inform the design and governance of these capabilities.

5. Discussion

Practical contribution. Regulation (EU) 2025/327 requires every Member State to allow the secondary use of electronic health data starting in March 2029, creating a future regulatory pathway for the advanced analytics function removed from the initial EDS deployment [7, 8].

Timed to anticipate this legislative deadline, the study offers three practical insights. First, it establishes an empirical checklist of governance conditions required for secure automated processing of health data, directly benefiting the Ministry of Health, AGENAS, and other Member States preparing similar systems. Second, it offers socio-technical design recommendations to help mitigate further protracted regulatory deadlocks, as seen during the iterative redesign period following the 2022 non-admission. Finally, it provides an empirical account of Italy's initial regulatory encounter with AI, offering lessons for other Member States as they develop their national health data infrastructures in preparation for the EHDS.

Theoretical contribution. This study builds on the cultivation-of-the-installed-base perspective [2] and applies it to the AI era. It extends the Introduction's argument that refusal can function as a resilience mechanism by showing how infrastructural resilience can also operate through regulatory and architectural gatekeeping: the ability to refuse, separate, and control when automated components are introduced. This finding aligns with and extends Mumford's socio-technical warnings about the risks of automated decision-making. Her concept of "dangerous decisions" helps frame the Garante's refusal of the 2022 analytics scheme not as a system failure, but as a deliberate form of infrastructural defence against large-scale automated processing of health data pursued without sufficiently transparent and participatory governance [87]. Refusal here operates at the level of specific configurations rather than at the level of AI itself: what was not admitted was a particular legal and architectural arrangement, subsequently redesigned rather than abandoned.

The study further identifies the architectural partitioning of AI capabilities within information infrastructures as a specific mechanism through which automated capabilities can be constrained, decoupled, or staged. The contribution is therefore problematising rather than gap-filling [47]: the case challenges the assumption that maintaining infrastructure and adapting to change are exclusively human activities, suggesting instead that resilience can also be produced through the governance of automated capabilities—including refusal, decoupling, and staging—across the three conceptual levels described in Section 3: epidemiological, operational, and clinical.

Studying an inactive part is intentional: resilience is mostly hidden when a system works well and only becomes clear when the system faces a disturbance and reacts. E-health research highlights an important analytical lens for studying these objects. Bygstad identifies two knowledge regimes [44].

- *Heavyweight IT* is managed by IT experts and built through careful planning and back-end software development.
- *Lightweight IT* is driven by users who require practical solutions and is often deployed through front-end applications by users and vendors themselves, without involving IT departments

Italy's two artificial intelligence trajectories align with Bygstad's two regimes: the analytics layer represents heavyweight AI integrated into the back-end core, while MIA serves as lightweight AI at the clinical edge. Bygstad's design principle states that the two regimes should be loosely coupled rather than tightly integrated, both technically (in terms of standards) and organisationally [44]. This principle helps explain the Italian arrangement: MIA's required separation from patient data is an example of loose coupling between lightweight and heavyweight regimes. The regulator establishes this separation in three ways: no technical access to the data, a certified guideline base replaces shared data flows, and AGENAS maintains separate organisational ownership. This paper's objective is closely linked to Bygstad's point that loose coupling helps units manage external changes and breakdowns and supports local adaptation. These are the same qualities that this study calls resilience. Because lightweight solutions can become heavyweight when fully integrated and managed under strict governance [44], this distinction shows that connecting MIA to patient data is not just a simple upgrade, but a bigger change that the health system must prepare for.

6. Conclusion

National health information infrastructures have remained resilient largely through human adaptability. This research-in-progress paper explores what happens to that resilience when part of this adaptive work is delegated to, or automated by, AI components. Italy provides a revelatory case: in one trajectory, AI-enabled analytics were proposed but later excluded from that legal and architectural configuration; in another, an AI assistant was admitted during its trial phase only under conditions that separated it from patients' personal data. The sequence does not describe an AI capability first excluded and then

automatically reintroduced. It describes a path that produced a regulatory and architectural design which distinguishes AI from the ordinary algorithmic functions supporting interoperability, and which makes any future admission — including secondary use under the European Health Data Space from 2029 — conditional rather than automatic.

Rather than filling a gap, this study problematises a central assumption in the information-infrastructures literature: that human actors are primarily responsible for maintenance, judgment, and trust-building. It advances a conceptualisation of resilience as a health system's capacity to adapt across epidemiological, operational, and clinical levels, understood as graceful extensibility and sustained adaptability rather than as rebound. This capacity is produced, in part, through the governance of admission: infrastructural resilience depends on the ability of institutions to subordinate the integration of AI components to defined legal bases, functional boundaries, technical and organisational safeguards, and human oversight.

The documentary phase is now complete. Next, we will conduct the main phase of expert interviews across different regions between late 2026 and early 2027, using the three short scenarios described in Section 3. We will then analyse the interview material alongside the completed documentary analysis and refine the emerging theoretical patterns. In 2027, we will consolidate the findings, incorporate additional publicly available evidence from the MIA trial and the evolving EHDS implementation, and develop the study's comparative and theoretical implications.

The study aims to deliver two contributions: conceptually, to illuminate how a health system admits, delimits, or declines to admit AI capabilities as expressions of infrastructural resilience; practically, to develop actionable design guidance for governing and implementing large-scale health data ecosystems under the evolving EHDS framework toward 2029.

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude and Grammarly to check grammar and spelling. Furthermore, the author(s) used NotebookLM to develop Figures 1, 2, 5, 6 and 7. After using these tools, the author(s) reviewed and edited the content as necessary and take(s) full responsibility for the content of the publication.

References

[1] S. L. Star, K. Ruhleder, Steps toward an ecology of infrastructure: Design and access for large information spaces, Information Systems Research 7 (1996) 111–134. doi:10.1287/isre.7.1.111.

[2] M. Aanestad, M. Grisot, O. Hanseth, P. Vassilakopoulou, Information infrastructures and the challenge of the installed base, in: Information Infrastructures within European Health Care: Working with the Installed Base, Springer, Cham, 2017, pp. 25–33. doi:10.1007/978-3-319-51020-0_3.

[3] Garante per la protezione dei dati personali, Parere sullo schema di decreto del Ministero della Salute sull'Ecosistema Dati Sanitari (EDS), 26 settembre 2024, doc. web n. 10062302. URL: https://www.garanteprivacy.it/home/docweb/-/docweb-display/docweb/10062302.

[4] M. G. Medici, M. Boni, Il nuovo decreto sull'Ecosistema Dati Sanitari: una rivoluzione digitale per la sanità italiana, Il Sole 24 Ore – NT+ Diritto, 3 marzo 2025. URL: https://ntplusdiritto.ilsole24ore.com/art/il-nuovo-decreto-sull-ecosistema-dati-sanitari-rivoluzione-digitale-la-sanita-italiana-AGtKgpHD.

[5] AGENAS, Piattaforma di Intelligenza Artificiale a supporto dell'assistenza primaria - Progetti strategici, n.d. URL: https://www.agenas.gov.it/agenzia-per-la-sanit%C3%A0-digitale/progetti-strategici/intelligenza-artificiale.

[6] AGENAS, Piattaforma di intelligenza artificiale: al via la collaborazione con le Regioni e le Province autonome per la sperimentazione delle attività, 2025. URL: https://www.agenas.gov.it/aree-tematiche/comunicazione/primo-piano/2665-piattaforma-di-intelligenza-artificiale-al-via-la-collaborazione-con-regioni-e-province-autonome-per-la-sperimentazione-delle-attivit%C3%A0.

[7] European Parliament and Council, Regulation (EU) 2025/327 of 11 February 2025 on the European Health Data Space, Official Journal of the European Union, 2025. URL: https://eur-lex.europa.eu/eli/reg/2025/327/oj/eng.

[8] European Commission, European Health Data Space Regulation (EHDS) - implementation timeline, n.d. URL: https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en.
[9] M. Aanestad, T. B. Jensen, Building nationwide information infrastructures in healthcare through modular implementation strategies, Journal of Strategic Information Systems 20 (2011) 161–176. doi:10.1016/j.jsis.2011.03.006.
[10] O. Hanseth, B. Bygstad, Flexible generification: ICT standardisation strategies and service innovation in health care, European Journal of Information Systems 24 (2015) 645–663.
[11] C. U. Ciborra, O. Hanseth, From tool to Gestell: Agendas for managing the information infrastructure, Information Technology & People 11 (1998) 305–327.
[12] S. L. Star, G. C. Bowker, How to infrastructure, in: L. Lievrouw, S. Livingstone (EDS.), Handbook of New Media, Sage, London, 2002, pp. 230–245.
[13] O. Henfridsson, B. Bygstad, The generative mechanisms of digital infrastructure evolution, MIS Quarterly 37 (2013) 907–931.
[14] O. Hanseth, E. Monteiro, M. Hatling, Developing information infrastructure: The tension between standardisation and flexibility, Science, Technology & Human Values 21 (1996) 407–426.
[15] O. Hanseth, K. Braa, Technology as traitor: Emergent SAP infrastructure in a global organisation, in: ICIS 1998 Proceedings, Association for Information Systems, 1998.
[16] S. L. Star, The ethnography of infrastructure, American Behavioural Scientist 43 (1999) 377–391.
[17] J. Vertesi, Seamful spaces: Heterogeneous infrastructures in interaction, Science, Technology & Human Values 39 (2014) 264–284.
[18] D. Tilson, K. Lyytinen, C. Sørensen, Research commentary - Digital infrastructures: The missing IS research agenda, Information Systems Research 21 (2010) 748–759.
[19] J. Braa, O. Hanseth, A. Heywood, W. Mohammed, V. Shaw, Developing health information systems in developing countries: The flexible standards strategy, MIS Quarterly 31 (2007) 381–402.
[20] P. N. Edwards, S. J. Jackson, M. K. Chalmers, G. C. Bowker, C. L. Borgman, D. Ribes, M. Burton, S. Calvert, Knowledge Infrastructures: Intellectual Frameworks and Research Challenges, Deep Blue, Ann Arbour, MI, 2013.
[21] S. J. Jackson, P. N. Edwards, G. C. Bowker, C. P. Knobel, Understanding infrastructure: History, heuristics and cyberinfrastructure policy, First Monday 12 (2007).
[22] T. P. Hughes, Networks of Power: Electrification in Western Society, 1880–1930, Johns Hopkins University Press, Baltimore, MD, 1993.
[23] E. Monteiro, N. Pollock, O. Hanseth, R. Williams, From artefacts to infrastructures, Computer Supported Cooperative Work (CSCW) 22 (2012) 575–607.
[24] J. Zittrain, The Future of the Internet - and How to Stop It, Yale University Press, New Haven, CT, 2008.
[25] J. Zittrain, The generative internet, Harvard Law Review 119 (2006) 1974–2040.
[26] V. Pipek, V. Wulf, Infrastructuring: Toward an integrated perspective on the design and use of information technology, Journal of the Association for Information Systems 10 (2009) 447–473.
[27] Ministero della Salute, Decreto 7 settembre 2023 - Fascicolo sanitario elettronico 2.0, 2023 (It.; disponibile su www.normattiva.it).
[28] Italian health data system: Current data interconnection and digital health ecosystem evolution, The Lancet Regional Health – Europe (2025) 101259. URL: https://www.thelancet.com/journals/lanepe/article/PIIS2666-7762(25)00051-1/fulltext.
[29] Y. Murgia, R. Gazzarata, M. Ciampi, M. Sicuranza, F. Cirillo, C. Esposito, N. Maggi, G. Balestra, L. Sacchi, M. Giacomini, The challenges of national health data ecosystems in feeding the European Health Data Space: The Italian example, Frontiers in Medicine 12 (2025) 1644719. doi:10.3389/fmed.2025.1644719.
[30] Ministero della Salute, Decreto 31 dicembre 2024, n. 334 - Istituzione dell'Ecosistema dei dati sanitari (EDS), Gazzetta Ufficiale, Serie Generale n. 53, 5 March 2025. URL: https://www.gazzettaufficiale.it/atto/serie_generale/caricaDettaglioAtto/originario?atto.dataPubblicazioneGazzetta=2025-03-05&atto.codiceRedazionale=25A01321&elenco30giorni=false.
[31] AGENAS, Ecosistema Dati Sanitari (EDS) - Progetti strategici, n.d. URL: https://www.agenas.gov.it/agenzia-per-la-sanit%C3%A0-digitale/progetti-strategici/EDS.
[32] Corte dei conti, Deliberazione n. 24/2025/G - Il Fascicolo Sanitario Elettronico (PNRR, M6C2), Sezione centrale di controllo sulla gestione delle Amministrazioni dello Stato, 2025. URL: https://www.corteconti.it/Download?id=4a19a61e-93d2-48ae-ae2f-d77eddfff9bd

[33] MIA, l'intelligenza artificiale per il medico di famiglia, Medico e Paziente, 2026. URL: https://medicoepaziente.it/2026/mia-intelligenza-artificiale-per-il-medico-di-famiglia/.
[34] R. Reale, E. Biasin, A. Scardovi, S. Toro, The design and implementation of a national AI platform for public healthcare in Italy: Implications for semantics and interoperability, arXiv preprint arXiv:2304.11893, 2023. doi:10.48550/arXiv.2304.11893.
[35] Garante per la protezione dei dati personali, Parere al Ministero della Salute sullo schema di decreto sull'Ecosistema Dati Sanitari (EDS), 22 agosto 2022, doc. web n. 9802752. URL: https://www.garanteprivacy.it/home/docweb/-/docweb-display/docweb/9802752.
[36] Garante per la protezione dei dati personali, Parere sullo schema di decreto sul Fascicolo Sanitario Elettronico 2.0, 8 giugno 2023, doc. web n. 9900433. URL: https://www.garanteprivacy.it/home/docweb/-/docweb-display/docweb/9900433.
[37] Svolta digitale in sanità: Big data e tutele individuali, Altalex, 4 giugno 2026. URL: https://www.altalex.com/documents/news/2026/06/04/svolta-digitale-sanita-big-data-tutele-individuali.
[38] Ministero della Salute, Sanità digitale: istituita la Cabina di regia nazionale per l'attuazione dello Spazio europeo dei dati sanitari, 2026. URL: https://www.salute.gov.it/new/it/news-e-media/notizie/sanita-digitale-istituita-la-cabina-di-regia-nazionale-lattuazione-dello.
[39] AI Act e Legge 132/2025: Cosa cambia per l'IA in sanità e per i dati clinici, Agenda Digitale, 2026. URL: https://www.agendadigitale.eu/sanita/ai-act-e-legge-132-2025-cosa-cambia-per-lia-in-sanita-e-per-i-dati-clinici/.
[40] G. Walsham, Interpretive case studies in IS research: Nature and method, European Journal of Information Systems 4 (1995) 74–81.
[41] A. Langley, Strategies for theorising from process data, Academy of Management Review 24 (1999) 691–710.
[42] R. K. Yin, Case Study Research and Applications: Design and Methods, 6th ed., Sage, Thousand Oaks, CA, 2018.
[43] N. Posteraro, Il fascicolo sanitario elettronico, in: V. Bontempi (Ed.), Lo Stato digitale nel Piano Nazionale di Ripresa e Resilienza, RomaTrE-Press, Roma, 2022, pp. 187–199. URL: https://romatrepress.uniroma3.it.
[44] B. Bygstad, Generative innovation: A comparison of lightweight and heavyweight IT, Journal of Information Technology 32 (2017) 180–193.
[45] Attacco ransomware all'ASL 1 Abruzzo e violazione dei dati: analisi del GDPR e della gestione del rischio informatico, ICT Security Magazine, 2026. URL: https://www.ictsecuritymagazine.com/articoli/attacco-ransomware-asl-1-abruzzo/.
[46] A Pescara nasce l'IoT-Cybersecurity Lab, il primo polo del Centro Italia di Rete8. URL: https://www.rete8.it/cronaca/123pescara-nasce-lot-cybersecurity-lab-il-primo-polo-del-centro-italia/.
[47] M. Alvesson, J. Sandberg, Generating research questions through problematization, Academy of Management Review 36 (2011) 247–271.
[48] P. Aspden, J. M. Corrigan, J. Wolcott, S. M. Erickson (Eds.), Components of a national health information infrastructure, in: Patient Safety: Achieving a New Standard for Care, National Academies Press, Washington, DC, 2004, Ch. 2.
[49] E. Klecun, A. Kankanhalli, Y. Zhou, Advancing understanding of scaling health information infrastructures: Learning from EHR initiatives in England, MIS Quarterly 49 (2025) 759–776. doi:10.25300/MISQ/2024/17307.
[50] E. Klecun, Y. Zhou, A. Kankanhalli, Y. H. Wee, R. Hibberd, The dynamics of institutional pressures and stakeholder behaviour in national electronic health record implementation: A tale of two countries, Journal of Information Technology 34 (2019) 292–332.
[51] C. Edwards, Information systems maintenance: An integrated perspective, MIS Quarterly 8 (1984) 237–256. doi:10.2307/249094.
[52] M. E. Kruk, M. Myers, S. T. Varpilah, B. T. Dahn, What is a resilient health system? Lessons from Ebola, The Lancet 385 (2015) 1910–1912.
[53] E. Hollnagel, J. Braithwaite, R. L. Wears (Eds.), Resilient Health Care, Ashgate, Farnham, 2013.
[54] S. Graham, N. Thrift, Out of order: Understanding repair and maintenance, Theory, Culture & Society 24 (2007) 1–25.
[55] S. J. Jackson, Rethinking repair, in: T. Gillespie, P. J. Boczkowski, K. A. Foot (Eds.), Media Technologies: Essays on Communication, Materiality, and Society, MIT Press, Cambridge, MA, 2014, pp. 221–239.
[56] S. Lebovitz, N. Levina, H. Lifshitz-Assaf, Is AI ground truth really true? The dangers of training and evaluating AI tools based on experts' know-what, MIS Quarterly 45 (2021) 1501–1525.

[57] E. van den Broek, A. Sergeeva, M. Huysman, When the machine meets the expert: An ethnography of developing AI for hiring, MIS Quarterly 45 (2021) 1557–1580.
[58] *J. O. Kephart, D. M. Chess, The vision of autonomic computing, Computer 36 (2003) 41–50. doi:10.1109/MC.2003.1160055.*
[59] P. N. Edwards, S. J. Jackson, G. C. Bowker, C. P. Knobel. Understanding infrastructure: dynamics, tensions, and design. Report of a workshop on "History & Theory of Infrastructure: Lessons for New Scientific Cyberinfrastructures". Ann Arbor: Deep Blue; 2007.URL: http://hdl.handle. net/2027.42/49353.
[60] O. Hanseth, K. Lyytinen, Design theory for dynamic complexity in information infrastructures: The case of building the internet, Journal of Information Technology 25 (2010) 1–19. doi:10.1057/jit.2009.19.
[61] M. Grisot, O. Hanseth, A. A. Thorseng, Innovation of, in, on infrastructures: Articulating the role of architecture in information infrastructure evolution, Journal of the Association for Information Systems 15 (2014) 197–219. doi:10.17705/1jais.00357.
[62] Y. Yoo, O. Henfridsson, K. Lyytinen, Research commentary — The new organising logic of digital innovation: An agenda for information systems research, Information Systems Research 21 (2010) 724–735. doi:10.1287/isre.1100.0322.
[63] B. Eaton, S. Elaluf-Calderwood, C. Sørensen, Y. Yoo, Distributed tuning of boundary resources: The case of Apple's iOS service system, MIS Quarterly 39 (2015) 217–243. doi:10.25300/MISQ/2015/39.1.10.
[64] B. Bygstad, E. Øvrelid, Architectural alignment of process innovation and digital infrastructure in a high-tech hospital, European Journal of Information Systems 29 (2020) 220–237.
[65] J. Wareham, P. B. Fox, J. L. Cano Giner, Technology ecosystem governance, Organisation Science 25 (2014) 1195–1215. doi:10.1287/orsc.2014.0895.
[66] T. P. Hughes, The evolution of large technical systems, in: W. E. Bijker, T. P. Hughes, T. Pinch (Eds.), The Social Construction of Technological Systems, MIT Press, Cambridge, MA, 1987, pp. 51–82.
[67] J. Porra, Colonial systems, Information Systems Research 10 (1999) 38–69.
[68] P. Weill, M. Broadbent, Leveraging the New Infrastructure: How Market Leaders Capitalise on Information Technology, Harvard Business School Press, Boston, MA, 1998.
[69] T. Kayworth, V. Sambamurthy, Facilitating localised exploitation of enterprise-wide integration in the use of IT infrastructures: The role of PC/LAN infrastructure standards, The Data Base for Advances in Information Systems 31 (2000) 54–80.
[70] P. A. Freeman, Is 'designing' cyberinfrastructure – or, even, defining it – possible?, First Monday 12 (2007).
[71] R. Kling, W. Scacchi, The web of computing: Computing technology as social organisation, in: Advances in Computers, Vol. 21, Academic Press, New York, 1982, pp. 3–87.
[72] R. Kling, Behind the terminal: The critical role of computing infrastructure in effective information systems' development and use, in: W. Cotterman, J. Senn (Eds.), Challenges and Strategies for Research in Systems Development, John Wiley, London, 1992, pp. 153–201.
[73] A. Zimmerman, A socio-technical framework for cyberinfrastructure design, in: Proceedings of the e-Social Science Conference, Ann Arbor, MI, 2007.
[74] P. Grindley, Standards, Strategy and Policy: Cases and Stories, Oxford University Press, New York, 1995.
[75] O. Hanseth, Information Technology as Infrastructure, PhD Thesis, University of Gothenburg, Gothenburg, 1996.
[76] P. A. David, S. Greenstein, The economics of compatible standards: An introduction to recent research, Economics of Innovation and New Technology 1 (1990) 3–41.
[77] H. A. Simon, The Sciences of the Artificial, 3rd ed., MIT Press, Cambridge, MA, 1996.
[78] European Parliament and Council, Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2 Directive), Official Journal of the European Union, 2022; transposed in Italy by d.lgs. 4 September 2024, n. 138.
[79] European Parliament and Council, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act), Official Journal of the European Union, 2024.
[80] European Parliament and Council, Regulation (EU) 2024/2847 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act), Official Journal of the European Union, 2024.
[81] European Parliament and Council, Regulation (EU) 2017/745 on medical devices and Regulation (EU) 2017/746 on in vitro diagnostic medical devices, Official Journal of the European Union, 2017

[82] Ministero della Salute, Dipartimento per la Trasformazione Digitale, Monitoraggio del Fascicolo Sanitario Elettronico 2.0 – Utilizzo, dati aggiornati al 31 maggio – 30 giugno 2026. URL: https://monitopen.fse.salute.gov.it/usage.
[83] Ministero della Salute, Decreto del 19 novembre 2025 — Disciplina del trattamento dei dati personali nell'ambito dell'infrastruttura della Piattaforma nazionale di telemedicina, Gazzetta Ufficiale, Serie Generale n. 301, 30 December 2025.
[84] H. K. Klein, M. D. Myers, A set of principles for conducting and evaluating interpretive field studies in information systems, MIS Quarterly 23 (1999) 67–93
[85] E. Mumford, Designing Human Systems for New Technology: The ETHICS Method, Manchester Business School, Manchester, 1983.
[86] E. Mumford, Effective Systems Design and Requirements Analysis: The ETHICS Method, Macmillan, London, 1995.
[87] E. Mumford, Dangerous Decisions: Problem Solving in Tomorrow's World, Kluwer Academic/Plenum Publishers, New York, 1999.
[88] European Parliament and Council, Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 as regards the timeline for the application of certain provisions (Digital Omnibus on AI), Official Journal of the European Union, 2026.
[89] A. Holmström, S. Andersson, J. Wenngren, Towards operationalising cyber resilience — a socio-technical analytical framework, CEUR Workshop Proceedings 4134 (2025). URL: https://ceur-ws.org/Vol-4134/paper7.pdf.
[90] E. Wynn, H. Vallo Hult, Qualitative and Critical Research in Information Systems and Human-Computer Interaction: Divergent and Convergent Paths, Foundations and Trends in Information Systems 3 (2019) 1–233. doi:10.1561/2900000014.
[91] Garante per la protezione dei dati personali, Decalogo per la realizzazione di servizi sanitari nazionali attraverso sistemi di intelligenza artificiale, settembre 2023.
[92] P. M. Bednar, C. Welch, Socio-Technical Perspectives on Smart Working: Creating Meaningful and Sustainable Systems, Information Systems Frontiers 22 (2020) 281–298. doi:10.1007/s10796-019-09921-1.
[93] P. Bednar, C. Welch, IS, Process and Organisational Change and their Relationships to Contextual Dependencies, in: Proceedings of the 13th European Conference on Information Systems (ECIS), ECIS Standing Committee, Regensburg, 2005, pp. 1466–1476.
[94] P. Bednar, C. Welch, A Double Helix Metaphor for Use and Usefulness in Informing Systems, Informing Science 10 (2007) 273–295. doi:10.28945/468.
[95] T. R. McEvoy, S. J. Kowalski, Deriving Cyber Security Risks from Human and Organisational Factors – A Socio-technical Approach, Complex Systems Informatics and Modelling Quarterly 18 (2019) 47–64.
[96] P. Bednar, C. Welch, M. Sadok, IS and Cybersecurity Practice: Avoiding Self-Sabotage, in: Proceedings of the 9th International Conference on Socio-Technical Perspectives in IS (STPIS 2023), CEUR Workshop Proceedings 3598, 2023, pp. 138–145.
[97] Legge 23 settembre 2025, n. 132, Disposizioni e deleghe al Governo in materia di intelligenza artificiale, Gazzetta Ufficiale, Serie Generale n. 223, 25 September 2025.
[98] E. B. Swanson, The dimensions of maintenance, in: Proceedings of the 2nd International Conference on Software Engineering (ICSE '76), San Francisco, 1976, pp. 492–497.
[99] E. Hollnagel, The four cornerstones of resilience engineering, in: C. P. Nemeth, E. Hollnagel, S. Dekker (Eds.), Resilience Engineering Perspectives, Volume 2: Preparation and Restoration, Ashgate, Farnham, 2009, pp. 117–133.
[100] D. D. Woods, Four concepts for resilience and the implications for the future of resilience engineering, Reliability Engineering & System Safety 141 (2015) 5–9. doi:10.1016/j.ress.2015.03.018.