Founda Health

White Paper

Beyond Consent

Why the Next Generation of Privacy Must Be Runtime-Governed

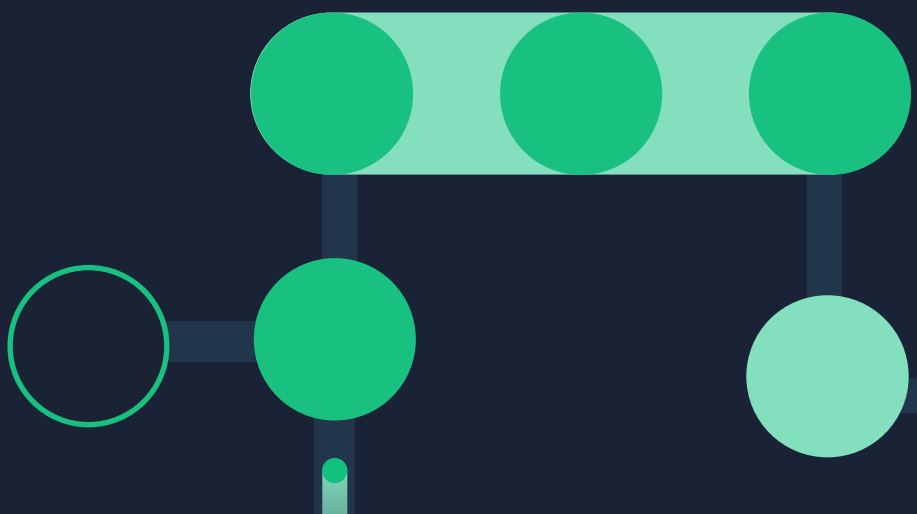

Privacy has been built as a body of law. The next decade will demand that it also become a layer of architecture: governance anchored to the person, carried with the data relationship, and evaluated continuously as that data moves between systems, organizations, and AI models.

Why intelligent health depends on trust we have not yet built

Picture the system working exactly as designed. A hospital wants to share data with a research partner under a signed contract. A governance committee has reviewed the request. An audit log will record every access.

Everyone involved is acting in good faith. And the answer is still no, because nobody in the chain can demonstrate, at the moment it matters, that the sharing is safe, purposeful, and authorized. That is the paradox at the center of healthcare data today: the institutions of trust exist, and the system still cannot move.

Healthcare is on the threshold of something genuinely transformative. The combination of always-available digital records, AI-assisted diagnosis, population health analytics, and personalized treatment pathways promises a future in which care is faster, more precise, and more proactive than anything the current system can deliver.

Not because clinicians will be replaced, but because technology will do what technology is good at, processing vast amounts of information at scale, so that clinicians can do what only humans can do well.

But this future has a precondition:

It requires data. Lots of it, flowing continuously, across many organizations and systems, for many purposes. A diagnostic AI model requires access to data it can process to provide value. A personalized care pathway requires longitudinal records that span decades and institutions. Population health management requires data from millions of individuals, shared across systems that were never designed to talk to each other.

And here is the problem opened above. We do not yet have the trust infrastructure to make that data flow legitimately. Today, most healthcare data sharing happens within narrow, pre-agreed boundaries: a clinician requesting prior records to treat a patient in front of them, a hospital sharing data with a specific research partner under a specific contract. Anything beyond those narrow cases runs into a wall of legal uncertainty, compliance caution, and institutional risk aversion. Information security officers say no.
Compliance officers say no. Not because the data sharing would be harmful, but because the governance mechanisms to demonstrate that it is safe, purposeful, and authorized simply do not exist in a form they can rely on.

The consequence is a healthcare system that cannot scale intelligently. We cannot afford to deliver personalized healthcare by scaling the workforce to meet every individual need. Technology, and the data it requires, is not optional. But technology without trust is not a solution, it is a liability. The pathway to intelligent health runs not around the trust problem but through it.

This article is about that problem. Not privacy as a compliance exercise. Not consent as a legal checkbox.

But the architecture of trust that intelligent healthcare requires, and why building it is one of the most important infrastructure challenges of the next decade. That architecture already has a precedent.

Cybersecurity faced this exact maturity gap two decades ago, and so did identity management more recently. Privacy is the discipline still catching up, and the rest of this article follows that path: how security and identity solved this problem, and what it means for privacy to solve it too.

The privacy problem this creates

For much of the digital age, privacy has been treated primarily as a legal problem, and understandably so.

Regulatory frameworks such as the GDPR [1] established important foundations: organizations must justify their processing, document lawful bases, report breaches, and give individuals rights over their information.

Compared with the largely unregulated data ecosystems that preceded it, this was significant progress.

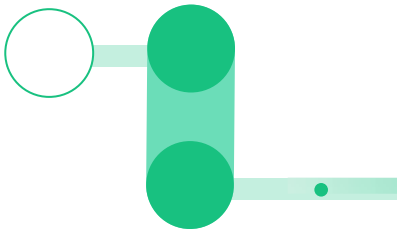

Yet despite these advances, something important remains unresolved. Most people still experience digital privacy as something happening around them rather than with them. A patient may originally disclose information to a provider for the purpose of direct care.

That information may then legitimately flow into wider shared care systems, regional interoperability exchanges, population health platforms, operational planning systems, research environments, AI-assisted diagnostic tooling, or external analytics providers. Each step along that path can be individually defensible: a legal basis exists, a contract covers it, a committee has signed off. Yet from the perspective of the individual whose data is traveling through all of it, the operational reality is often opaque. They rarely have real visibility into where their data is, who is using it, how it is being used, whether the original purposes still apply, or how downstream uses evolve over time.

The problem is not that regulation failed. It is that regulation alone was never going to be sufficient. Modern privacy frameworks evolved strongly as a legal governance mechanism, but only weakly as an operational one. That gap is what this article addresses, and it matters more than it first appears.

The precedent in identity governance

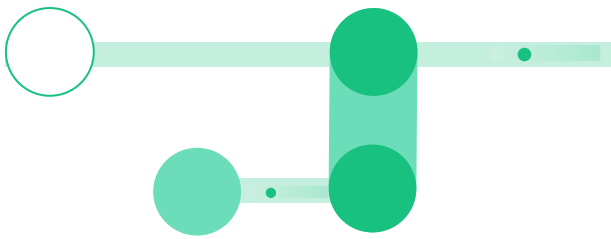

A useful comparison can be found in Identity Governance and Administration. Access rules are not merely documented and forgotten. Permissions are monitored and audited in real time, violations are flagged automatically, and access can be revoked on the spot.

Governance is not simply policy, it is an active, live layer woven through the architecture itself.

Privacy regulation rarely functions this way today. Once data has crossed an organizational boundary, governance frequently becomes fragmented. The receiving organization may inherit contractual obligations, but the governance context surrounding the data often weakens technically as it propagates through systems, copies, transformations, and derived datasets.

Privacy becomes primarily a matter of institutional process rather than continuous runtime governance.

Early signs the shift is possible

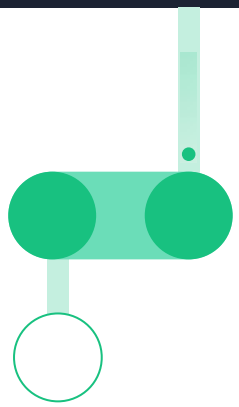

Early exceptions are instructive. In the Netherlands, the Mitz initiative represents one of the first meaningful steps toward operationalizing subject-expressed consent within a live healthcare data ecosystem. Through integration with DigiD, the national digital identity infrastructure, Mitz enables individuals to express and manage their own consent preferences through a government-backed system, which healthcare providers then actively evaluate at the point of processing decisions.

This is not merely a legal notice mechanism. It is the beginning of a runtime consent layer, one in which the subject of the data participates operationally, not just retrospectively through terms of service.

The limitation is that this consent layer is today country-specific and not yet standardized. A subject's consent is expressed and evaluated within one national system. It does not yet travel with the data across the wider ecosystem, and today, data rarely crosses a national border at all unless a patient carries it there themselves. That gap, portable, standards-based consent that moves with the data between processors, is precisely the missing piece, and it is exactly what the European Health Data Space is designed to close (more on this below). Mitz demonstrates that the transition is feasible. The question is how to generalize it.

AI changes the shape of the problem

This challenge becomes even more pronounced in the era of artificial intelligence. Traditional privacy models were built around a relatively simple mental model of data processing: data moves from one organization to another for a defined purpose

AI disrupts this assumption. Personal data may now contribute probabilistically toward model training, influence latent representations, shape future outputs, or participate in derived inferences that become increasingly difficult to trace back to their origins.

This raises a governance question that modern privacy architecture must be capable of answering, even where a definitive answer remains contested:

> Once data has helped shape a model, can governance still follow it, or does the trail effectively end at training time? Should governance frameworks evolve in step with processing purposes as they develop over time?

These are not arguments against AI in healthcare. On the contrary, AI-assisted care represents one of the most significant opportunities to address capacity constraints and improve clinical outcomes across European health systems. They are architectural requirements: governance infrastructure must be capable of maintaining meaningful observability and accountability across these new processing patterns, rather than assuming that traditional data-transfer models are sufficient for a fundamentally different computational paradigm.

Governance that travels with the relationship

So far, this has been diagnosis. Here is where the argument turns to the answer.

Modern privacy systems still largely assume governance can remain external to the data itself. Policies sit in documents. Consent sits in databases. Contracts sit between institutions. Governance becomes attached to organizations rather than to the relationship the data describes. This model increasingly struggles under the weight of distributed healthcare ecosystems, interoperable platforms, and AI-driven processing environments.

The next generation of privacy may therefore require a different architectural assumption.

It is tempting to express this as governance that travels with the data, but that framing is slightly misleading, and the distinction matters. Governance does not need to ride inside the payload to travel with it. What needs to be continuous is the relationship between the person, the parties they have authorized, and the purposes they have agreed to.

If that relationship is expressed once, anchored to the person, and made verifiable at runtime, governance follows the data wherever it goes without having to be physically attached to it.

This does not imply a simplistic decentralization model, nor does it require the elimination of institutional trust. Hospitals, research institutions, public health authorities, and technology providers will continue to play essential roles in society. Nor is this an argument against GDPR.

On the contrary, frameworks such as GDPR established many of the legal foundations required for accountable digital processing. The challenge is that legal governance alone may no longer be sufficient for increasingly dynamic and distributed digital ecosystems. What may now be required is an operational governance layer capable of continuously representing and evaluating the permissions, obligations, and provenance associated with a person's data as it moves between systems, anchored to the person, rather than re-derived by each institution that happens to hold a copy.

For those of us building healthcare data availability infrastructure, this reframes the task. The interoperability layer is not only where data is made available. It is the natural place to make governance observable as that data moves.

A consent credential, and who should carry it

It is worth being precise about what this governance layer would actually look like in practice, and what it would not. One instinct is to imagine a privacy envelope that wraps every individual piece of data as it moves between systems, carrying governance metadata alongside the payload.

That framing has intuitive appeal, but it overcomplicates the mechanics. Attaching cryptographic governance metadata to every data object at transfer time is operationally unwieldy and architecturally unnecessary.

A cleaner model separates two distinct concerns: identity and consent. They look similar, but they have opposite trust topologies, and recognizing that is what unlocks the design.

The first is identity. Europe is already building this layer. Under eIDAS 2.0, every EU member state is required to issue citizens with a digital identity credential, held in a European Digital Identity Wallet, by 2027 [2]. This credential establishes who the person is, anchored to a state-issued identifier, and is designed to be presented to any organization that needs to verify identity, without the issuing authority being involved in every transaction.

Think of it as owning your own car but driving on roads owned by the state: more individual control than before, built on infrastructure that requires trusted institutional foundations. Identity is the part that genuinely needs an external root. A self-asserted claim to be a particular person is worthless to a hospital unless something the hospital already trusts binds that claim to a real, legal individual.

That is exactly what state-issued identity provides, and it is why this model does not attempt to decentralize identity itself.

Consent is the mirror image. Its authority does not originate from the state or the institution, it originates from the person. Verifying who someone is requires an outside authority to vouch for the claim. Verifying that someone consents does not: the person is the only authority the claim needs.

Consent is therefore the one part of this architecture that can legitimately be self-attested by the subject, provided it is bound to an externally anchored identity. That asymmetry lets us keep the institutional anchor exactly where it is irreducible, identity, and remove it exactly where it is not, consent.

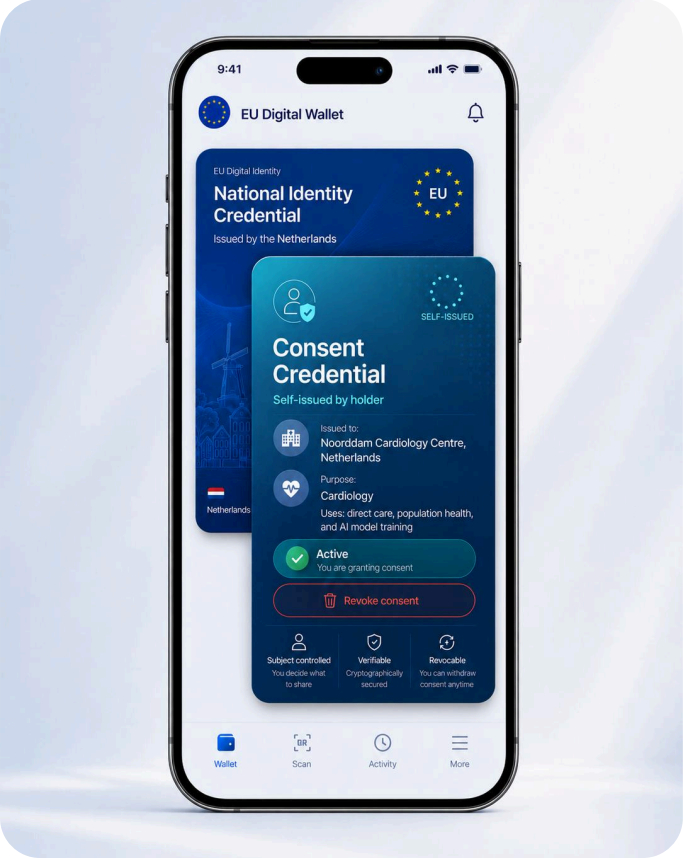

This points to a different design from the one most consent systems adopt today, and it is worth being explicit about the inversion involved.

In today's models, the burden of consent sits with the data custodian, the organization that stores the data. Before disclosing, the custodian must establish that the subject has consented.

But the custodian rarely holds the live relationship with the patient, and cannot anticipate every future purpose for which the data might one day be legitimately requested.

This produces one of two unsatisfactory outcomes. Either consent collapses into something so broad as to be meaningless, in the Netherlands this has produced what is now called generic consent, along the lines of "I permit this organization to share my data if needed," which the patient cannot truly understand because the conditions are left undefined. Or the system tries to enumerate permissions in advance through a central service and a fixed policy matrix.

Mitz takes the second path: it lifts the administrative burden off the custodian by centralizing consent, but because it too cannot predict future use, it must ask the patient a large set of abstract, up-front questions, on the order of two dozen permutations of "do you allow a provider of type A to access your data," and it cannot express anything that falls outside that matrix.

The deeper problem is that the custodian is being asked to govern a relationship it does not hold. The live, purpose-bearing relationship belongs to the requestor, the provider the patient has actually come to for care, not to the custodian sitting on records from years ago. So invert it. Rather than the custodian checking consent, the requestor presents evidence of consent as part of the request. The subject is the issuer of that evidence, the requesting provider becomes its holder, the custodian becomes a verifier.

In practice, consent is captured at the moment it is most meaningful, and most answerable. When a patient registers with a new provider, whether in person or through an online portal, they are asked to consent to that specific, concrete engagement. Using their digital identity wallet on their phone, they sign a consent credential: an expression of what they consent to, with whom, for what purpose, for how long, and where it is valid. Because the question is tied to a real, present need rather than a hypothetical future one, the patient can actually reason about it, the opposite of the abstract, pre-emptive questions a central matrix is forced to ask.

During registration, the provider presents a QR code, either on screen or in their portal. The patient scans it with their wallet. The wallet retrieves a credential template published by the provider: a structured document that describes the consent being requested, pre-populated with the provider's own verified identity, the specific purposes for which they will process the patient's data, and the duration and scope of the relationship. Think of it as a partially completed form. The provider has specified what they need, the patient reviews it, adds their own identity binding (a cryptographic link to their state-issued national identity credential that establishes who they are beyond doubt), confirms the terms, and signs the credential using their wallet. The result is a consent credential issued by the patient themselves, held in their wallet, and verifiable by any party who receives it.

This inversion, where the patient issues rather than receives, is unfamiliar but important. It shifts the authority over the consent relationship to the person it belongs to, without requiring that person to become a technical expert. The wallet handles the cryptography, the patient makes the human decision.

It is worth being honest about where the standards for this specific flow currently stand. The building blocks exist: the OpenID for Verifiable Credentials family of specifications covers both how credentials are offered and how they are presented, and the European Digital Identity Wallet architecture is being built on these foundations.

But the precise profile for a patient-issued consent credential of this kind, including the credential offer and self-issuance flow described here, has not yet been formally standardized. This is engineering and standards work that remains to be done. The architecture is sound, the standardization is in progress.

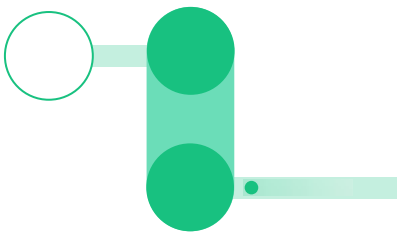

Most of these dimensions are closed and finite, the parties, the duration, the jurisdiction, and standardize trivially. The harder dimension is purpose: what, exactly, the data may be used for. Purpose can be arbitrarily rich, and the instinct to standardize all of it is precisely what drives systems back toward an unmanageable matrix. The resolution is to separate two layers within the credential. A coarse, standardized purpose, a small, closed set of codes of the kind healthcare interoperability already defines, is the part any verifier can reason over. Beneath it sits the actual agreement between the subject and the requesting provider, which can be as specific as the relationship requires and which no third party needs to interpret.
The thin, evaluable layer is standardized, the rich detail stays a matter between the two parties.

The transfer right then follows as a consequence, not a separate grant. If a provider is authorized by the patient to process their data for a given purpose, they are thereby authorized to obtain that data from whoever holds it. When the provider approaches a custodian for prior records, they present the credential, selectively disclosing only the facts the custodian needs to see: that the patient issued it, to this provider, valid for this period, for this coarse purpose, together with a pointer to check its current status.

The custodian learns what it needs to make a disclosure decision and nothing more, the detailed agreement between patient and provider is never exposed to it.

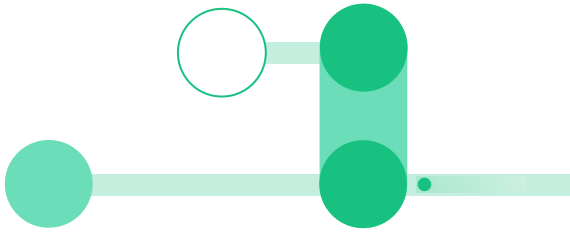

Critically, the custodian is not reduced to a rubber stamp. It remains a constrained verifier. It checks three things: that the credential is authentic and was issued by the patient, that it has not been revoked, and that the data being requested falls within the coarse purpose the patient authorized. A credential issued for cardiac care does not entitle a requestor to a patient's entire history.

The transfer right then follows as a consequence, not a separate grant. If a provider is authorized by the patient to process their data for a given purpose, they are thereby authorized to obtain that data from whoever holds it.

The custodian releases what fits the stated purpose and withholds the rest. This is what keeps disclosure proportionate. It also means the custodian's responsibility does not disappear so much as transform, from the open-ended, unpredictable task of administering consent for unknowable future purposes, into a bounded, evidence-backed proportionality check against a credential the requestor brings with them.

That is a real reduction in burden, but it is a reduction, not an elimination, and it is precisely this residual check that makes governed data a meaningful claim rather than a marketing one. It also places a real demand on the semantic layer: deciding what fits a purpose requires data classified well enough that the question can be answered, which is the same standards-based foundation interoperability already depends on.

Revocation is what turns this from a point-in-time event into a living relationship. When a patient issues a consent credential, they do not hand over control permanently. They can withdraw consent at any time, and when they do, that withdrawal needs to be detectable by any organization currently relying on that credential, without requiring the patient to contact each of them individually. This is solved through a revocation mechanism built into the credential itself.

Each consent credential contains a pointer to a status list, a publicly readable, machine-checkable record of whether that credential is still valid. Technically, this uses a standard called the W3C Bitstring Status List [3]: a compact format in which each credential is assigned a position in a long sequence of bits, each of which can be flipped from valid to revoked. The list is published at a known address by the infrastructure the patient's member state provides as part of the digital wallet framework. Any organization holding a credential can check that address at any time to see whether the patient has since withdrawn consent.

From the patient's perspective, this is simple. They open their wallet, find the relevant consent credential, and revoke it. The wallet updates the status list. The next time the processor checks, they will find consent has been withdrawn. The patient does not need to contact anyone. They do not need to know which systems hold copies of their data. The revocation propagates automatically through the infrastructure.

This is what makes governance continuous rather than historical.

A processor cannot simply check consent once at first access and assume it remains valid indefinitely. As compliance monitoring matures, processors will be expected to check revocation status periodically, and a credential that has been revoked is no longer a valid basis for processing. The patient's change of mind has operational force, not just legal expression. The status list infrastructure is provided as part of the member state's digital wallet framework, the individual does not need to run their own server or manage their own endpoint. This is the thin, shared layer the model requires: not a central authority that owns consent, but a lightweight, federated mechanism that lets consent remain verifiable without remaining static. Governance becomes a continuously observable property of the relationship rather than a decision frozen at first contact.

Importantly, this does not imply perfect technical enforcement. Human intent remains difficult to model computationally. Organizations acting in bad faith may still attempt misuse, just as malicious insiders still exist within modern cybersecurity systems. Yet the objective is not perfection.

The objective is the reduction of invisible trust assumptions through increased observability, traceability, and operational accountability.

In many respects, this mirrors the evolution that already occurred in cybersecurity. Modern security architectures do not assume that prevention alone is sufficient. They combine prevention, observability, runtime evaluation, telemetry, and governance into continuous operational systems. Privacy governance may now be approaching a similar transition.

This is, in effect, an extension of what a data availability platform already does. Making the right data available in the right place is the first problem. Carrying the governance of that data with it across processors, anchored to the person, in a standardized way rather than one confined to a single country, is the next.

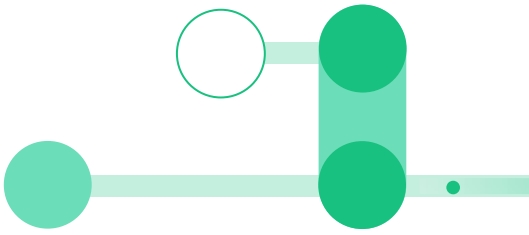

EHDS is already creating the demand

This transition is already visible in emerging regulatory architecture. The European Health Data Space, EU regulation 2025/327 [4], establishes a mandatory framework for cross-border health data exchange across EU member states, with progressive implementation through 2031.

Critically, EHDS does not merely expand the scope of data sharing. It creates governance requirements around secondary use, data access conditions, and individual rights within highly distributed data ecosystems.

EHDS creates the regulatory demand for exactly the kind of consent credential layer described here, and eIDAS 2.0 [2] creates the identity infrastructure on which that layer can be built. The two programs, proceeding in parallel across EU member states, are establishing the technical and legal foundations for a genuinely portable governance model, one in which a citizen's consent is anchored to their state-issued identity, expressed through a standardized credential, and verifiable by any authorized processor across the European health data ecosystem.

The technical architecture required to meet EHDS obligations at scale will, by necessity, push organizations toward machine-readable governance metadata, traceable provenance, and runtime evaluation of processing permissions, precisely because the complexity and distribution of the mandated exchanges cannot be managed through institutional process alone.

Notably, EHDS secondary use operates on an opt-out basis rather than active individual consent, reflecting the practical limitations of current governance models. This is understandable as a near-term policy choice, but it should be recognized for what it is: a reliance on the very institutional trust assumption this article argues we must move beyond. Opt-out works only for as long as that trust holds. If people lose confidence in who is using their data downstream, and for what, the rational response is to opt out, and at scale, mass opt-out would starve precisely the longitudinal, secondary-use data that intelligent, personalized care depends on.

A more mature governance architecture, in which a person can see and steer their own data relationships, is not merely a refinement of opt-out. It is what makes the trust that opt-out quietly assumes sustainable in the first place.

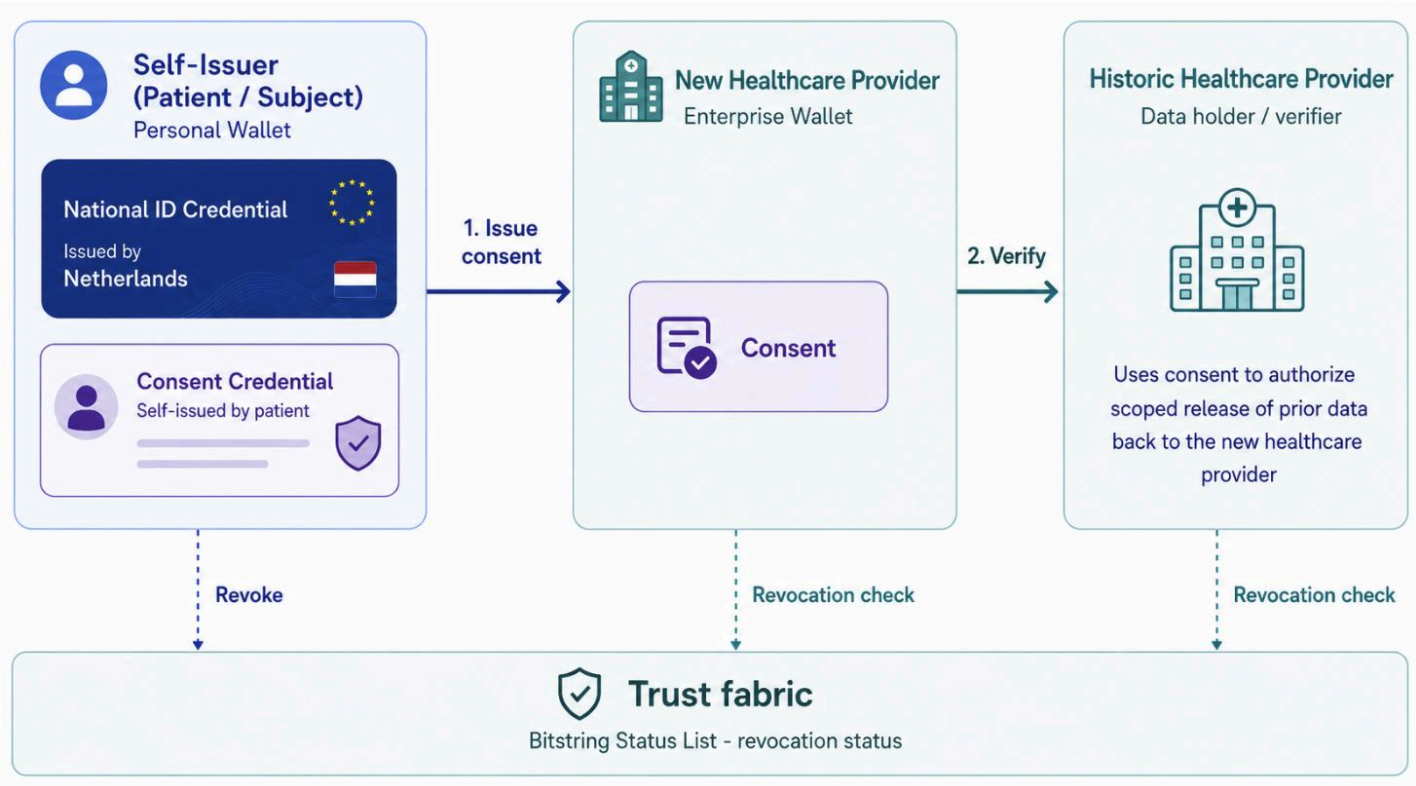

The person, not the institution

There is also a deeper philosophical shift emerging beneath these technical considerations. Modern digital systems implicitly model institutions as primary actors and individuals as subordinate entities represented within institutional systems. Hospitals create patient identifiers. Platforms create accounts. Organizations define the boundaries of identity relationships. Yet institutions are transient compared with the continuity of the person. Healthcare providers change. Technology vendors change. Governments change. Systems are replaced repeatedly over decades. The person remains.

This becomes especially significant in longitudinal healthcare contexts such as lifelong care records, population health management, and AI-assisted preventative medicine, where personal data may outlive individual systems, organizations, and even generations of technology architecture. The natural person, rather than the institution, may therefore need to become the primary anchor point for future privacy governance architectures. Not because institutions become unimportant, but because the continuity of governance ultimately belongs to the subject whose life the data describes.

The next generation of privacy

The privacy debates of the last decade focused heavily on regulation, consent, and compliance. The next decade may focus increasingly on operational governance, provenance, and runtime accountability. As healthcare ecosystems become more interoperable, AI-driven, and distributed, the limitations of static privacy models will become increasingly difficult to ignore. Legal governance remains essential, but legal governance alone cannot continuously observe, evaluate, and operationalize data relationships across highly dynamic digital environments.

Cybersecurity evolved from static policy toward runtime enforcement. Identity governance evolved from manual administration toward continuous operational control. Privacy governance may now be approaching the same architectural threshold.

The likely path is not pure decentralization. Governments will continue to issue foundational identities because societies need widely trusted anchors. But the emerging model, state-issued identity combined with subject-controlled consent credentials, verified at runtime by any processor in the chain, is a meaningful advance on what exists today. It is decentralized enough to give individuals genuine agency over their data relationships, and standardized enough to function across the complexity of a European health data ecosystem.

The next generation of privacy will not be defined solely by stronger regulation. It will be defined by the emergence of runtime-governed privacy architectures capable of continuously representing, evaluating, and operationalizing the interests of the individual within increasingly intelligent and interconnected systems.

At Founda, we see this as the next layer of data availability. The first task is to make the right information reach the right point of care. The next is to ensure the governance of that information travels with the person whose data it is, observable, accountable, and anchored to the individual whose life the data describes. .

That is the direction we are building toward: not privacy as a promise made once and filed away, but governance carried in the infrastructure itself.

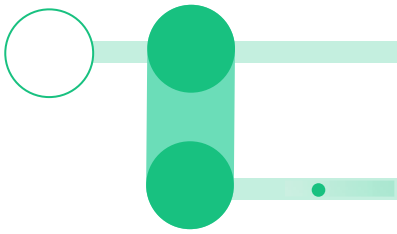

References

1. Complete guide to GDPR compliance. https://gdpr.eu/
2. eIDAS 2.0 | Updates, Compliance. https://www.european-digital-identity-regulation.com/
3. Bitstring Status List v1.0: Privacy-preserving status information for Verifiable Credentials. W3C. https://www.w3.org/TR/vc-bitstring-status-list/
4. European Health Data Space Regulation (EHDS). European Commission. https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en

Founda Health

About the author

Rob Walker

Chief Technology Officer

Rob has spent over 16 years leading technology teams, with more than a decade building national-scale healthcare technology. He led the engineering behind eRedbook - the UK's digital personal child health record - and the Digital Staff Passport, and contributed to national NHS programmes including NHS Login. His expertise spans personal health records, healthcare interoperability standards and identity and access management, and he has advised organisations including the World Bank on digital identity. As Chief Technology Officer at Founda Health, he leads the technology strategy for making healthcare data available wherever and whenever it is needed.

About Founda Health

Founda Health enables secure, standards-based availability of healthcare data between organizations and systems. The Founda Platform helps healthcare networks connect existing infrastructure, support national interoperability initiatives, and prepare for future European requirements around data sharing and consent.