Consent Architecture: DKMS as the Structurally Sound Foundation for EHDS Article 71

Vereign AG

2026-04-28

Consent Architecture: DKMS as the Structurally Sound Foundation for EHDS Article 71

Consent Architecture: DKMS as the Structurally Sound Foundation for EHDS Article 71

Vereign AG — April 2026

1. Executive Summary

European healthcare data law has crossed a threshold. Opt-out is no longer aspirational policy — under EHDS Article 71 it is a **legal right** of the natural person, enforceable across every Member State that processes identifiable health data for secondary use. Layered onto GDPR Articles 7(3) and 17(2), eIDAS 2.0 Article 5a, the Swiss FADP and HFG, and the German PDSG, five independent regimes converge on a single set of engineering invariants: **person-scoped identifiers, verifiable timestamped opt-out, reversibility without re-identification, cross-controller propagation, and unlinkability**. Centralised PKI and federated identity satisfy at most one or two; X.509 satisfies none of them at the citizen layer. **Decentralized Key Management (DKMS) is the first structurally sound foundation for honoring opt-out end-to-end *in cross-organisational, post-disclosure, regulator-watched conditions*.** The thesis has explicit boundaries — derivatives that left the system before withdrawal, immutable backups, and single dishonest counterparties remain out of architectural reach. The 2025–2026 settlements docket (Kaiser USD 47.5M, Sutter USD 21.5M, parallel filings against BJC, Northwell, Catholic Health, Aspirus, and SSM Health) is the leading indicator: every centralised incident shares one feature — the user had no architectural agency.

2. Why Now — Regulatory Convergence in 24 Months

For two decades, healthcare-IT consent management treated opt-out as a UX problem. From 2024 onward, five independent legal regimes converged on the same engineering pattern, leaving vendors with months — not years — to demonstrate structural compliance.

EHDS Article 71 — Regulation (EU) 2025/327. Adopted 11 February 2025; published in the Official Journal of the European Union (OJ L 2025/327) on 5 March 2025; in force from 26 March 2025. Article 71(1) gives every natural person the right to opt out of secondary processing of identifiable electronic health data. Article

71(2) prohibits new data permits over identifiable data once opt-out is exercised, while preserving permits issued before the exercise. **Article 71(8)** forbids data holders from acquiring extra identification data solely to honour the opt-out — the engineering pivot that forecloses the obvious "central register of opted-out IDs" pattern.

GDPR Articles 7(3) and 17(2) — Regulation (EU) 2016/679. Article 7(3) requires that withdrawal of consent be "as easy to withdraw as to give." Article 17(2) imposes a propagation duty: a controller honouring an erasure request must take "reasonable steps, including technical measures" to inform downstream controllers. Article 9(2)(a) raises the threshold for health data to *explicit* consent; Article 17(3)(c)/(d) provides public-health and research carve-outs that make literal erasure structurally unreachable for most clinical records — which is precisely why EHDS framed the right as **forward suppression** rather than retrospective deletion.

eIDAS 2.0 — Regulation (EU) 2024/1183. Amends Regulation (EU) 910/2014 to establish the European Digital Identity Wallet (EUDIW). Article 5a(14) mandates "full user control" over wallet operations. Article 5a(16)(a) prohibits cross-context tracking of wallet usage; Article 5a(16)(b) mandates **unlinkability** between credential presentations to different verifiers. Together, 5a(16)(a)–(b) are the operational form of EHDS Article 71(8): an architecture in which the operator can correlate a citizen across verifiers fails the unlinkability test by definition.

Swiss FADP and HFG. The revised Federal Act on Data Protection (FADP / nDSG, SR 235.1, in force 1 September 2023) requires explicit consent for sensitive data at Article 6(6) and an absolute withdrawal right at Article 6(7), with Article 32 deletion rights subject to retention obligations. The Human Research Act (HFG / HRA, SR 810.30) at Article 7(2) and Article 17 grants revocation of consent and a right to dissent for further use of health data. The incoming EGDG (federal council message, 5 November 2025) inverts the EPD model from opt-in to opt-out with go-live ~2028–2030; the BDG (DigiSanté secondary-use legal basis) is in drafting. Cross-cantonal propagation makes the Swiss case structurally harder than the EU one.

German PDSG. The Patientendaten-Schutz-Gesetz operationalises the German ePA (electronic patient record) opt-out architecture, in live operation from January 2025 across roughly 73 million statutory-insurance beneficiaries. The PDSG codifies forward-suppression mechanics — citizens declining the ePA must not be processed under the opt-in legal basis — across a population larger than most EU Member States combined. National implementation pressure on consent architecture is therefore not theoretical; it is operational in the largest single-payer healthcare estate in Europe.

The TEHDAS2 Joint Action draft guideline (September 2025) confirms that the technical mechanism of opt-out propagation is **not** specified by EU law itself and is deferred to Member State implementation. This is the architectural soft underbelly: each state could pick its own approach, but only a subject-controlled-keying approach satisfies Article 71(8) without smuggling in a forbidden re-identification register. Vendors that ship a compliant architecture in one Member State will sell across all twenty-seven; vendors who solve only the local UX will face a different architectural conversation in every transposition.

Each one of these five regimes alone is survivable. **All five together, in a 24-month window, are the architectural forcing function.** Healthcare-IT vendors are now confronted with a market in which "consent management" must compose end-to-end, across borders, across controllers, across time — and the centralised

registry pattern that has carried the industry for two decades is foreclosed by the primary law of the largest economic bloc on the planet.

---

3. The Five Engineering Invariants

The five regimes converge on a small, precise set of engineering requirements. They are not aspirational — they are testable predicates against which any consent architecture can be evaluated.

3.1 Person-Scoped Identifiers

The data subject must control the identifier under which their permits, authorizations, and consent state are referenced. This is the *prerequisite* for every other invariant: if the operator owns the identifier, the operator owns the opt-out. EHDS Article 71(8) explicitly forecloses operator-side re-identification registers; eIDAS 2.0 Article 5a(14)'s "full user control" mandate names the same property at the wallet layer. Without a subject-scoped identifier, propagation, reversibility, and unlinkability all collapse into operator policy.

3.2 Verifiable Timestamped Opt-Out

For each data permit, the system must know the precise moment of issuance and the precise moment of withdrawal — and a third party (regulator, court, downstream controller) must be able to verify both without trusting the operator's logs. EHDS Article 71(2) draws a sharp temporal boundary: permits before the opt-out remain valid; permits after it are unlawful. This is unenforceable without a tamper-evident, third-party-verifiable timeline of the subject's consent state.

3.3 Reversibility Without Re-Identification

Recital 54 of the EHDS Regulation confirms that opt-out is reversible (citizens can opt back in) and free-form (no minimum duration). Article 71(8) prohibits the data holder from acquiring extra identification data solely to honour the opt-out. The two clauses read together create a tight constraint: the system must support flipping consent state on and off, on the citizen's timing, **without** the data holder building a parallel re-identification register to track who has flipped which way. Any architecture that requires "just keep a list of opted-out IDs" violates 71(8) directly.

3.4 Cross-Controller Propagation

GDPR Article 17(2) requires the controller honouring erasure to take "reasonable steps, including technical measures" to inform downstream controllers holding copies, replications, or links. EHDS Article 71 layered onto multi-controller permit chains escalates this: the opt-out signal must reach every Health Data Access Body, every research consortium, every sub-processor that received data under a previous permit. Propagation cannot be a paperwork register — it must be a verifiable artefact downstream parties can re-check.

3.5 Unlinkability

eIDAS 2.0 Article 5a(16)(b) mandates unlinkability between presentations of the same credential to different verifiers. The operational consequence: an architecture in which any single party (even an authorized intermediary) sees every transaction across verifiers fails the test by construction. This is the property that excludes federated identity providers from the citizen layer of EHDS-compliant systems: the IdP, by design, sees every login and every credential presentation. Federation can satisfy propagation, but only by violating unlinkability.

These five invariants form an interlocking system. **X.509 PKI satisfies zero of them at the citizen layer. Federated IdPs satisfy propagation only by violating unlinkability. DKMS satisfies all five.** That is the structural-soundness claim, and the rest of this document examines it.

---

4. The Article 71(8) Trap

The simplest implementation pattern for opt-out — the one a healthcare-IT vendor will draft in the first hour of architectural conversation — is **a central registry of opted-out citizen IDs**. Every data holder queries the registry before processing; if the citizen's ID is on the list, processing stops. Easy. Familiar. Auditable.

Article 71(8) forecloses this pattern entirely.

The clause prohibits data holders from acquiring extra identification data solely to honour opt-out. **The "list of opted-out IDs" is, by construction, exactly such a register.** It exists for no purpose other than re-identifying citizens for the negative case — to determine that *this citizen has opted out and therefore cannot be processed*. The list itself is the regulatory violation it is meant to prevent.

The trap is recursive. Any attempt to anonymise the registry (hash the IDs, pseudonymise them, keep them in a separate trust domain) re-introduces the original problem: the data holder still needs *some* identifier to test against the list, and that identifier must be derivable from the data the holder already has. The hash is not anonymous to the holder who has the input; the pseudonym is not anonymous to the holder who has the linkage. Recital 54 reinforces this by mandating that the opt-out be reversible and free-form — meaning the registry must support flipping, which in turn requires that the registry track *which* citizen has flipped *when*, which is exactly the re-identification record 71(8) prohibits.

The implication for healthcare-IT vendors is sharp. **Any "consent management module" whose architecture rests on a central authoritative list of opted-out subjects is foreclosed by EU primary law.** This is not a UX design choice; it is the difference between an architecture that can comply with EHDS and one that cannot. Generic consent-management products built on the registry pattern — whether SaaS, on-prem, or hybrid — will face this constraint at every Member State transposition. The structural answer must place the identifier under the citizen's control, so that the data holder never has to maintain a re-identification register at all.

---

5. The DKMS Thesis

> **Decentralized Key Management is the first structurally sound foundation for honoring opt-out end-to-end *in cross-organisational, post-disclosure, regulator-watched conditions*.**

DKMS — built on KERI, an open Trust over IP Foundation specification — places the trust root with the citizen rather than with a centralised authority. The citizen (or their wallet agent) holds an Autonomic IDentifier (AID) that is self-certifying: it is derived from a public key under the citizen's control, and its authenticity does not depend on any registrar. All authorizations granted to data holders, processors, and downstream controllers chain back to that AID through a Key Event Log (KEL) — an append-only, third-party-verifiable record of the citizen's keying state, including key rotations.

Withdrawal in this model is not a request submitted to an operator. It is a **key rotation the citizen performs unilaterally**. The rotation is published to the KEL; from that moment on, any downstream party holding credentials anchored to the pre-rotation key state must re-authenticate against the citizen's current state to continue processing. Existing permits issued before the rotation continue to operate where the underlying legal basis allows (matching EHDS Article 71(2)); new permits cannot be issued under the old key state, because the old state is no longer the citizen's controlling state. The citizen has architecturally — not just procedurally — withdrawn future authorization, and any downstream party that ignores the published rotation produces signatures that are auditable as stale.

Mapping this against the five invariants:

Invariant	DKMS mechanism
Person-scoped identifiers	The AID is the subject's self-certifying identifier; no operator-side index required
Verifiable timestamped opt-out	KEL contains tamper-evident, timestamped key events
Reversibility without re-identification	Opt-back-in = key rotation re-establishing authorization; no parallel register
Cross-controller propagation	KEL is publishable; verifiers must re-check across trust domains
Unlinkability	Per-context AIDs and selective-disclosure credentials avoid cross-verifier correlation

Compare this to existing options:

X.509 PKI binds identity to a CA-issued certificate. The citizen does not hold the trust root; the CA does. Revocation is operator-driven (CRL/OCSP) and soft-fails by default. Cross-context unlinkability is structurally absent — every certificate presentation is linkable through the certificate's serial number, issuer DN, and subject DN. X.509 satisfies *zero* of the five invariants at the citizen layer.

Federated IAM (OIDC, SAML). The Identity Provider is by design the cross-context tracker that eIDAS 2.0 Article 5a(16)(b) prohibits. Federation can satisfy cross-controller propagation (a single logout flushes downstream sessions), but only by violating unlinkability — every login flows through the IdP, which sees every relying party the citizen visits. The architecture can satisfy at most three invariants and only at the cost of the fifth.

Centralised KMS (AWS KMS, Azure Key Vault, GCP Cloud KMS). Excellent for single-tenant workloads. NIST SP 800-88 §2.5 cryptographic erasure via Customer-Managed Keys is a regulator-recognised terminal-erasure primitive (EDPB Guidelines 01/2025; ENISA 2019 §4.2). But once data leaves the tenancy boundary — into a research consortium, a downstream controller, a permitted secondary-use recipient — the operator's CMK destruction does nothing to the copy. Cross-controller propagation is unsolved. Schrems II / EDPB Recommendations 01/2020 Use Case 3 escalates this further: the data importer must not hold the keys, by definition.

DKMS does not replace these stacks — it occupies the precise architectural slot they cannot fill. Centralised KMS for tenancy-scoped erasure, X.509 for TLS-class revocation, federated IAM for session-scoped consent — and DKMS for the cross-organisational, post-disclosure layer where EHDS Article 71 actually lives.

---

6. Where Centralised Architectures Fail

The five invariants are testable. Each major centralised consent architecture fails in a structurally specific way.

6.1 X.509 PKI Fails All Five Invariants at the Citizen Layer

X.509 binds identity to a Certification Authority. The CA generates or co-signs the citizen's certificate, owns the issuance decision, and controls revocation. Person-scoped identifiers fail at the root — the CA's signature is the trust source, not the citizen's. Timestamped opt-out is operator-driven (the citizen submits a revocation request; the CA decides to honour it and when to publish), and the revocation propagates only as fast as CRL refresh windows or OCSP responder availability — both of which soft-fail in browsers by default. Reversibility without re-identification fails because re-issuance requires the CA to re-validate the subject's identity, which in healthcare contexts requires the CA to re-acquire the very identification data EHDS Article 71(8) prohibits the data holder from holding. Cross-controller propagation depends on every relying party correctly checking CRLs/OCSP, which the DigiNotar and Comodo incidents proved is operationally fragile. Unlinkability is structurally impossible: every X.509 certificate carries a serial number and issuer DN that link every presentation to the same citizen.

6.2 Federated IAM (OIDC/SAML) Fails Unlinkability by Definition

OIDC and SAML federate authentication through an Identity Provider that brokers the relationship between the citizen and the relying party. Every login the citizen performs against any verifier transits the IdP. The IdP knows which citizen accessed which service when. This is not a configuration choice or a policy violation — it

is the architectural function of the IdP. eIDAS 2.0 Article 5a(16)(b) mandates unlinkability between presentations of the same credential to different verifiers; the federated IdP is, by construction, the single point where every presentation is correlated. Federation can deliver real cross-controller propagation (Single Logout, OAuth Token Revocation, OIDC Back-Channel Logout) — but the cost is the violation of the unlinkability invariant for every credential presentation the IdP brokers.

6.3 Centralised KMS Cannot Reach Cross-Controller Propagation

AWS KMS, Azure Key Vault, and GCP Cloud KMS provide auditable key custody, FIPS 140-2 L3 / 140-3 backends, scheduled key destruction (`ScheduleKeyDeletion`, soft-delete + purge, `DESTROY_SCHEDULED`), and envelope encryption with DEK/KEK separation. Within a single tenancy, this is a fully adequate GDPR Article 17 erasure primitive — NIST SP 800-88 §2.5 explicitly recognises Cryptographic Erase as a Purge-level sanitisation technique, and EDPB Guidelines 01/2025 treat key custody as outcome-determinative for whether data is anonymous to a given recipient. The architecture works **inside the tenancy boundary**.

It does not cross the boundary. Once the data has been lawfully copied to a research consortium, a permitted downstream controller, or a sub-processor in another trust domain, the tenant's KMS destruction does nothing to the copy. GDPR Article 17(2) propagation becomes a contractual obligation — paperwork, not architecture. EHDS Article 71's permit-chain structure makes this gap structural rather than incidental: every permit is potentially the start of a multi-controller chain, and centralised CMK has no means of reaching past the first link.

---

7. Where DKMS Does NOT Deliver — Steel-Manned

The thesis is structural, not absolute. Honest framing makes the argument defensible; overreach destroys it. The DKMS claim has four explicit boundaries.

7.1 Derivatives — EDPB Opinion 28/2024

Once data has been transformed into a derivative — analytics aggregates, ML model weights, statistical extracts, regulatory reports already filed — key destruction on the source ciphertext does nothing to the derivative. EDPB Opinion 28/2024 on data protection aspects of AI models (adopted 17 December 2024) explicitly confirms that erasure obligations do not propagate cleanly to model weights once training has occurred. DKMS does no better than centralised CMK here. If a permitted research consortium has trained a model on data the citizen later opts out of, the model weights remain. The architectural contribution is auditability of the derivation chain, not retroactive unwinding of the derivation.

7.2 Backups Already Taken

No live-key revocation architecture, DKMS or centralised, can recall data already replicated to immutable backup media. Offline backups taken before the withdrawal sit outside the live keyspace. The pragmatic response is documented, scheduled key destruction on a known retention window — which ICO, CNIL, and BfDI guidance all endorse as sufficient for backup-residual personal data. DKMS aligns with this discipline; it does not replace it.

7.3 Single Dishonest Counterparty

In an N-party network with one party that ignores the published key rotation and continues to act on stale credentials, DKMS provides **forensic detection** — the stale signatures are auditable past the published rotation event — but not **prevention**. A fully non-cooperative downstream processor that exfiltrates plaintext and stores it outside the keyed envelope is not within architectural reach of any key-management system. DKMS strengthens the regulator's evidentiary position; it does not fix the underlying non-cooperation. Compared to a contractually well-governed centralised CMK estate, DKMS strengthens detection but does not eliminate the need for contract, regulator, or court enforcement against bad actors.

7.4 Single-Tenant Workloads

For workloads that live entirely within one trust domain — internal enterprise data under one controller, single-tenant SaaS with tenant-scoped CMK, public-web TLS under CA/Browser Forum 24-hour revocation SLOs — NIST SP 800-88 §2.5 cryptographic erasure via centralised CMK is fully adequate. EDPB Guidelines 01/2025 and ENISA 2019 §4.2 both recognise key destruction as terminal erasure inside the tenancy boundary. **Adding DKMS to these workloads imposes operational burden for no architectural benefit.** The honest answer is that for single-tenant scenarios, the centralised stack works.

7.5 The Decisive Boundary

DKMS dominates when **all four** of the following conditions hold:

1. Data crosses **at least two trust domains**.
2. Consent must **survive post-disclosure** (i.e., processing continues across the boundary).
3. The regulator treats **key custody as outcome-determinative** (Schrems II / EDPB Recommendations 01/2020 Use Case 3 / EHDS Article 71(8)).
4. Parties compose DKMS **end-to-end** (a verifier that ignores the KEL nullifies the chain).

Outside those four conditions, centralised CMK is the rational choice, and saying so makes the DKMS thesis stronger — not weaker — because it locates the claim precisely where the evidence holds. The claim is not that DKMS solves cookie banners or single-tenant SaaS deletion. The claim is that DKMS is the right structural choice for *exactly* the cases that matter for cross-cantonal, cross-organisational, regulator-watched, post-disclosure healthcare consent.

8. Production Proof

Vereign operates the DKMS architecture in production today, with explicit caveats on what is shipped versus what is architectural.

SEAL is the production communication layer: encrypted swarm delivery for outbound communication, with per-message keys and recipient-controlled decryption agency. SEAL has been carrying **800,000+ verified messages every month** for institutional users — the canonical production-scale anchor for the architecture. Each message exercises the same key-event substrate that underlies the consent thesis: subject-controlled keying, third-party-verifiable timestamps, and per-context unlinkability between recipients.

Stargate — the AID-anchored authorization layer where data permits chain to the citizen's KEL — enters early production from June 2026 with HIN as the first operational deployment. Stargate is the consent surface for the thesis: authorizations granted to data holders are anchored to the citizen's AID; key rotation by the citizen propagates through the KEL; downstream verifiers re-check. The HIN deployment is a soft launch in scope; the **mass roll-out narrative is reserved for post-summer 2026**, after the early-production cohort has accumulated operational evidence under live regulatory conditions.

HIN bridges allow withdrawal to interoperate with X.509-using counterparties via the KERI-anchored S/MIME bridge documented in the email-security knowledgebase. This is the pragmatic acknowledgement that DKMS will compose with — not replace overnight — the existing X.509 healthcare estate.

FHIR-over-Stargate is architectural — production deployment depends on hospital onboarding; earliest plausible September 2026. The architectural story is genuinely strong: FHIR resources flowing through Stargate inherit the AID-anchored authorization model, so resource-level consent state lives where the citizen controls it. But the engineering of hospital adoption — clinical-system integration, regulatory sign-off, operational governance — is a multi-quarter undertaking. We do not claim FHIR-over-Stargate as shipped. We claim it as the deliberate next step, with a known earliest date, and we expect operational deployment to lag the architectural readiness by a measurable interval. Conflating the two is the failure mode this document is written to avoid.

The production track record matters because it is what separates this thesis from a slide deck. Engineering organisations encounter the centralisation-vs-decentralisation argument routinely; what they encounter rarely is decentralisation operating at scale, under live institutional regulation, in a regulated industry. The verified-messaging substrate has been carrying real institutional traffic for long enough to expose the operational corners of DKMS — key rotation discipline, cross-organisation discovery, audit-log retention, recipient agency — and those corners have been refined in production rather than in theory. When Stargate enters early production from June 2026, it will run on the same substrate. The architectural argument is therefore not a forecast; it is an extrapolation from observed operations at scale.

9. Five Incidents That Prove the Structural Pattern

Five incidents — three from 2026, two foundational older cases — illustrate why centralised architectures fail under regulatory pressure and how the failure mode is architectural rather than operational.

9.1 Kaiser Foundation Health Plan — USD 47.5M Settlement (2026)

Kaiser Foundation Health Plan reached a USD 47.5 million class-action settlement over patient-portal pixel-tracking, with final approval entered on 30 April 2026 (Class Action Center, Kaiser Foundation Health Plan settlement docket). Plaintiffs alleged that Meta and Google tracking pixels embedded in Kaiser's authenticated patient portal transmitted protected health information to third-party advertising platforms — including diagnoses, appointment types, and medication queries — alongside identifiers sufficient to re-link the data to specific patients. The architectural failure is precise: patient consent existed at the registration layer of the portal, but the data flow was determined by the embedded SDKs, not by the patient's stated preference. Withdrawing consent did nothing to the pixel events already transmitted, and the patient had no architectural means to detect or revoke the third-party flows.

9.2 Sutter Health — USD 21.5M Settlement (2026)

Sutter Health reached a USD 21.5 million settlement in the same docket pattern, with final approval on 27 February 2026. Parallel settlements have been entered or are pending against BJC HealthCare, Northwell Health, Catholic Health, Aspirus, and SSM Health on essentially identical facts. The 2025–2026 settlement docket as a whole is the leading indicator: every major US health system that embedded third-party tracking pixels in patient-facing surfaces is now paying for the architectural absence of per-patient cryptographic boundaries on the data flow. Consent was a checkbox; the flow was decided elsewhere.

9.3 NHS National Data Opt-Out (NDOO) — Non-Retroactivity, 2018–2021

NHS England consolidated the legacy Type 1 and Type 2 opt-outs into the National Data Opt-Out (NDOO) from 2018, with statutory force. Programme documentation and ICO commentary subsequently confirmed that the opt-out was demonstrably **non-retroactive**: data already extracted from GP systems before the citizen registered the opt-out remained in research datasets and could continue to be processed. NHS Digital had no technical infrastructure to recall already-extracted records. The successor GPDPR (General Practice Data for Planning and Research) programme was scheduled to begin 1 July 2021 and was indefinitely delayed in June 2021 specifically because there was no technical means to delete already-extracted GP data when a Type 1 opt-out was registered after the fact. As of 2026 the extraction has still not begun. The architectural lesson: when the operator owns the data path, "patient veto" has no enforcement primitive — and the policy mechanism collapses under its own forensic weight.

9.4 SingHealth — 1.5M Records Compromised (2018)

The Singapore Committee of Inquiry's Public Report into the Cyber Attack on Singapore Health Services (10 January 2019) found that 1.5 million SingHealth patient records had been compromised, including the personal health record of the Prime Minister of Singapore. The technical finding is the architectural one: there was **no compartmentalisation** in the SingHealth electronic medical records system. Once an attacker reached the Allscripts Sunrise application via a privileged service account, all patient records were accessible. Patient consent was a policy artefact at the access layer; the data layer had **no per-patient cryptographic boundary**. The breach surface was precisely the absence of the boundary.

9.5 DigiNotar Certificate Authority Compromise (2011)

The Fox-IT investigation published as the "Black Tulip" report (13 August 2012) found that all eight DigiNotar CA servers had been compromised. 531+ rogue certificates had been issued, including a wildcard `*.google.com` certificate used to mount a man-in-the-middle attack against approximately 300,000 Iranian Gmail users. The remediation — distrusting DigiNotar — was made by Mozilla, Microsoft, and Google unilaterally and propagated to end users through root-store updates. The Dutch government's own Staat der Nederlanden – G2 chain was overlooked in the initial removal and required follow-up. End users had **zero agency** in the trust decision at any stage: not in the issuance, not in the detection, not in the remediation. They learned about the compromise via their own services breaking.

Every centralised incident shares one feature: the user had no architectural agency. Kaiser and Sutter patients did not control the pixel flow. NHS patients could not recall extracted data. SingHealth patients had no per-record cryptographic boundary. DigiNotar end users could not vote on trust decisions. The pattern is structural, and it persists across five decades of incident evidence regardless of regulator, sector, or jurisdiction.

---

10. Conclusion and How to Engage

Opt-out has crossed from policy aspiration to enforceable legal right under EHDS Article 71, GDPR 7(3) and 17(2), eIDAS 2.0 5a(14)/(16)(b), Swiss FADP and HFG, and the German PDSG. Five engineering invariants follow directly: person-scoped identifiers, verifiable timestamped opt-out, reversibility without re-identification, cross-controller propagation, and unlinkability. X.509 PKI satisfies zero at the citizen layer; federated IAM satisfies propagation only by violating unlinkability; centralised KMS satisfies the single-tenant case but cannot reach across trust domains. **DKMS is the first structurally sound foundation for honoring opt-out end-to-end in cross-organisational, post-disclosure, regulator-watched conditions** — with explicit boundaries around derivatives, backups, and single dishonest counterparties. The 2025–2026 settlements docket (Kaiser, Sutter, BJC, Northwell, Catholic Health, Aspirus, SSM Health) is the live evidence that every centralised architecture under this kind of regulatory pressure shares one feature: the user had no architectural agency.

If you are responsible for consent architecture in a healthcare-IT estate facing EHDS transposition, Swiss EGDG / BDG implementation, or German PDSG operations, Vereign offers a 30-minute consent architecture review. We will map your current consent surface against the five invariants, identify the Article 71(8) traps in your roadmap, and locate the precise scope where DKMS is — and is not — the right architectural answer. **Contact: contact@vereign.com.**

Vereign AG — Dammstrasse 16, 6300 Zug, Switzerland — UID CHE-240.299.384 — LEI 50670056G9BYC736YR76