openEHR and HL7® FHIR® in Europe

Evidence, Architecture and Procurement Implications

Assessment

Document 3 · Dr Kai U. Heitmann · Release 1.54.1 – 7 September 2026

The full argument, its derivation and the recommendation it supports.

Written for: Interoperability and standards specialists, and reviewers who need to examine how the recommendation was derived.

For the decision alone, see Document 1. For the executive argument and tender questions, see Document 2. Case material and the graded source ledger are in Document 4. The German profile-corpus analysis is in Document 5. Document 6 states what would have to change for the next version of this assessment to say more.

About this document

Internal section references point to this document. References to Appendices A to E and G point to the Evidence Annex (document 4). References to Appendix D-10 and profile-corpus counts point to the Technical Report (document 5). Appendix F is the glossary of technical terms and abbreviations.

Character and scope. This is an evidence-based technical assessment that ends in an explicitly identified procurement recommendation. It separates source-supported findings, analytical inference and authorial judgement. The normative premise is stated in the section "Recommendation and limits" so that it can be examined directly.

The geographical scope is Europe, where the European Health Data Space (EHDS), MyHealth@EU and national implementation rules apply and where openEHR has its strongest documented deployment base. Findings on regulatory anchoring and relative adoption are therefore region-specific. Architectural findings on model boundaries, transformation obligations and governance are more widely transferable. Examples outside Europe are used only where they inform feasibility.

The document has two parts. Part I provides a compact analytical orientation. Part II presents the technical derivation and source criticism.

Contents

Part I — Analytical Orientation

Purpose and scope

Interoperability is more than this assessment. Exchanging data in a standard format, or storing them in one, is only a part of interoperability. The dimensions are conventionally distinguished as technical connectivity, syntactic conformance, semantic consistency, process integration, organisational responsibility and governance. They are interdependent, and conformance to a data format establishes only one link in the chain. That is the point set out in full in § 1.5.

The OECD's 2026 cross-country assessment independently reinforces this wider framing. It distinguishes technical, semantic, legal, regulatory and organisational dimensions and identifies governance, trust, adoption, workforce capability and data quality as conditions of effective interoperability. Its analysis supports the boundary drawn here: data standards and model governance are necessary foundations, but they do not by themselves establish workflow integration, organisational adoption or beneficial outcomes (Document 4, Appendix E-10).

Within those dimensions this assessment examines the mechanisms that the choice and governance of clinical data models materially affect. Among them are profiling depth, terminology binding and the strength of a constraint. They are the means by which semantic interoperability becomes probable and, more importantly, checkable.

The mechanisms examined here fall into three groups.

- The first concerns the record: which model holds the **authoritative clinical content** and where it is first written, and how many **separately governed model boundaries** a record crosses.
- The second concerns the constraints: how tightly profiles and archetypes constrain what may be recorded, and how terminology is bound and at what **strength**.
- The third group concerns the governance: how transformations are **validated**, how conformance is tested and by whom, and how models, versions and export paths are governed over the **life of the data**. The consequences examined are those for **portability, procurement and patient safety**.

The conclusions concern the contribution of data architecture and standards governance to interoperability, and not the interoperability of a health system as a whole. Clinical workflow integration, organisational change, user-interface usability and health outcomes depend on applications, configuration, professional practice and institutional arrangements at least as much as on data standards. This document does not evaluate them.

This assessment examines how **openEHR and HL7® FHIR®** affect architecture and procurement in the European regulatory context. It does not ask which standard is universally better, but where clinical data are first written, which model remains authoritative, how many independently governed model boundaries the data must cross and what evidence supports the resulting choice.

FHIR is explicitly anchored in the European requirements for data exchange. openEHR has no comparable EU-level mandate, but it is used in several European clinical-data platforms. Regulatory status establishes planning relevance, not technical or clinical superiority.

Recent Swiss, Welsh, Australian, Austrian, Spanish and industry documents agree on FHIR as the exchange standard but differ on the persistence question and on the role assigned to openEHR. The Spanish document is the newest and reaches furthest: a ministry-commissioned Delphi consensus recommends openEHR as the most suitable standard for the national digital medical record, with FHIR at the exchange layer — a recommendation to the ministry, not an adopted strategy, and by its own words not evidence-based (Appendix E-12).

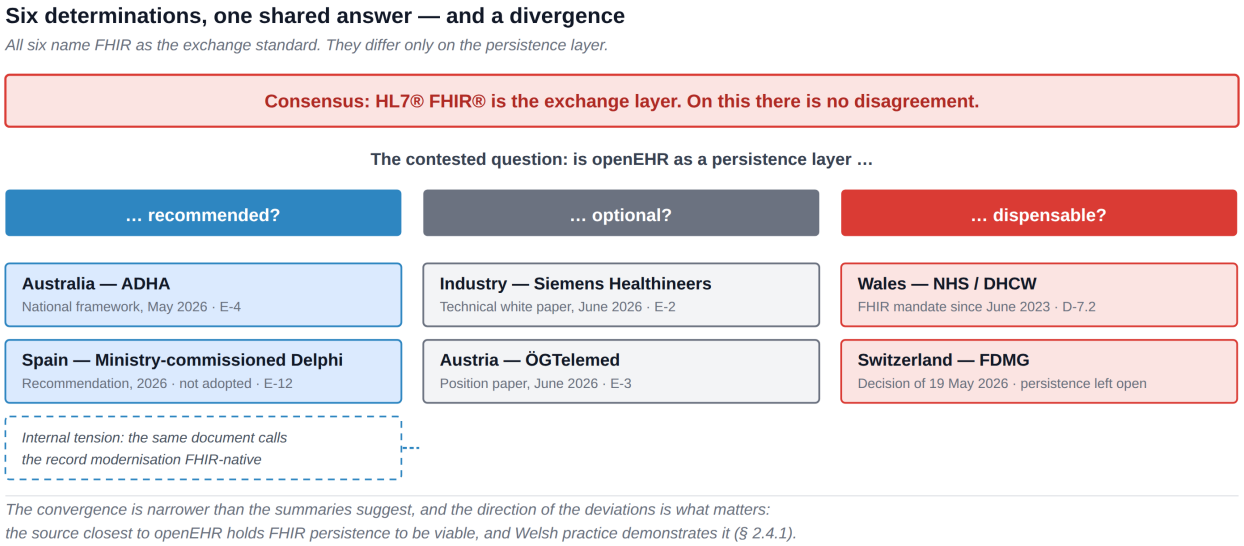

Figure 1 — Six determinations, 2023–2026. All of them name FHIR as the exchange standard, and they scatter across the whole range on the persistence question; one recommends openEHR for the record itself, and one points in both directions at once (§ 2.4.1; Appendix E-3, E-12).

Findings

1. The operational choice concerns five arrangements

The relevant distinction is not simply between two standards. It is between five ways of combining source and storage models:

	CLINICAL DATA ARE FIRST WRITTEN IN	AUTHORITATIVE CLINICAL MODEL	FHIR EXCHANGE	MINIMUM GOVERNED CONVERSIONS
A	Proprietary application model	Proprietary model	Generated when required	**1**
B	Proprietary application model	Proprietary model; the FHIR repository is a durable replica	Exposed from the FHIR replica	**1**
C	FHIR-native application	FHIR repository	Direct use of the authoritative FHIR representation	**0**, if compatible with the required exchange profile
D	Proprietary application model	Proprietary model; the openEHR repository is a durable replica	Transformed from the openEHR replica	**2**
E	openEHR-native application	openEHR repository	Transformed from openEHR	**1**

The count concerns boundaries between separately governed clinical models, not software layers, caches or transport steps. A boundary exists when the same clinical meaning must remain correct in two independently specified and versioned representations.

Arrangements A and E each require one transformation before FHIR exchange, and for a receiving system they are indistinguishable at that boundary. The case for openEHR therefore cannot rest on the exchange boundary but must rest on a demonstrated advantage in modelling depth, clinical governance, reuse or long-term stability.

Because European exchange requirements use FHIR, an openEHR record must cross at least one model boundary before exchange. A FHIR record can avoid that boundary. This asymmetry follows from the regulatory context and does not establish technical superiority.

2. Transformations create continuing obligations

Each governed model boundary requires mapping, terminology alignment, testing, clinical validation and change management. Tooling can reduce implementation effort but cannot remove responsibility for the semantic relationship between the models.

The number of transformations is an architectural indicator, not a measure of cost or safety. The actual risk depends on the content, direction, frequency, automation, terminology bindings, model stability and operational controls. One poorly governed transformation may be more hazardous than two well-controlled ones.

Every mapping path should therefore have named clinical responsibility, representative test data, acceptance criteria, regression testing and documented handling of content that cannot be mapped without loss. This requirement is especially important for medication, allergy and other safety-critical information.

3. Feasibility is demonstrated, comparative effectiveness is not

Persistent FHIR repositories and operational openEHR platforms demonstrate that both approaches are feasible. Large FHIR repositories show that FHIR can serve as a durable target model at national and institutional scale. They do not show that comprehensive primary documentation is already written directly as FHIR resources. In the cases examined here, data generally originate in primary systems and are transformed into the repository.

No head-to-head study identified for this assessment compares equivalent FHIR-native and openEHR-plus-FHIR architectures for clinical outcomes or full lifecycle cost. Published evidence also does not establish that one or two transformations produce different rates of patient harm.

FHIR has a clear lead in regulatory adoption, implementation breadth, community size and engagement, tooling and workforce availability. These are material procurement considerations. They are not evidence of greater clinical effectiveness.

4. Conformance alone does not produce interoperability

Applications may conform to the same standard while using incompatible profiles, extensions, terminology bindings or value sets. Only carefully governed applications can interoperate across a transformation boundary.

A standards decision must therefore be accompanied by profile governance, terminology services, conformance testing, version management and operational monitoring. Choosing FHIR or openEHR without this regime leaves the main implementation problem unresolved.

5. Vendor neutrality is a contractual and governance property

Neither standard prevents vendor dependency by itself. Such lock-in can arise through the repository product, cloud platform, proprietary source system, implementation-specific profiles or the model repository needed to interpret stored data.

An openEHR instance may remain dependent on the versions of the archetypes and operational templates under which it was created. A FHIR implementation may depend on local extensions, search parameters and platform services. Procurement must therefore secure long-term access to data, models, profiles, mappings, terminology bindings and the rights and tools required to use them after contract termination.

Recommendation and limits

> For genuinely new infrastructure, FHIR should be the common storage and exchange model. An additional openEHR layer is justified only where a demonstrated domain benefit outweighs the permanent cost and risk of maintaining another model transformation. This is a conditional procurement judgement, not a finding that FHIR is clinically or technically superior.

Here, *genuinely new* infrastructure means an environment in which the organisation can still choose the authoritative clinical model and is not constrained by an established repository, legacy application models or prior contractual commitments (the boundary is set out in § 1.6.3).

The recommendation may not apply to:

- **Existing systems:** Replacing a functioning architecture may cost more and create more risk than improving its interfaces and governance.
- **Secondary platforms:** Research, analytics and cross-organisational platforms may justify a separate canonical model because their purpose differs from primary documentation. The resulting secondary representation and its maintenance obligations should be made explicit.
- **Dedicated or specialised environments:** Established openEHR capabilities or a demonstrated need for its modelling and governance approach may outweigh the additional FHIR boundary.

The decision should be reconsidered if comparative evidence shows lower lifecycle cost, better model reuse, greater clinical completeness or better long-term stability for an openEHR-plus-FHIR architecture, or if FHIR-internal complexity removes the assumed advantage. A material change in European data exchange requirements would also alter its regulatory basis.

Procurement consequences

A tender should require explicit answers to five questions:

1. **Where is the authoritative clinical record?** Identify the model in which each clinical fact is first committed and from which other representations can be reconstructed.
2. **How is every transformation clinically validated?** Require representative sets of clinical test data, named responsibility, regression testing and controlled treatment of unmappable content. Demand evidence that changes in source models, profiles and terminology bindings do not go unnoticed.
3. **How are profiles and terminologies governed?** Require enforceable constraints, versioned value sets, explicit bindings and automated validation. Formal conformance without sufficient profiling does not ensure usable interoperability.
4. **Can the organisation leave the contract with both data and meaning intact?** Secure export and continued use of data, profiles, templates, mappings and terminology assets. Vendor neutrality depends on governance, intellectual-property rights, skills and exit provisions, not on the standard alone.
5. **What is the ten-year lifecycle obligation?** Include modelling, mapping, validation, upgrades, migration, workforce requirements and supplier change. Initial implementation cost alone does not carry the architectural choice.

Part II derives these findings, examines the counterarguments and records the limits of the underlying sources.

Part II — Technical Analysis

This part provides the full derivation of the findings summarised in Part I, the source criticism and the appendices. It starts with implications on a strategic and political level and is subsequently addressed to a specialist audience in medical informatics, architecture and standardisation.

1. Strategic and Political Level

This section establishes what the recommendation rests on that lies outside the standards themselves:

- what European law settles and what it leaves open (§ 1.1),
- what each side can fairly claim for itself (§ 1.2, § 1.4),
- where dependency actually arises (§ 1.3),
- why passing a conformance test is not the same as being interoperable (§ 1.5),
- what one further governed model boundary costs, in money and in patient safety (§ 1.6), and
- which points remain contested, rather than resolving them (§ 1.7).

The order is the order of the argument, not the order of weight. The section runs from the frame to the mechanism, and its heaviest part comes last. Read in this order if time is short:

1. **§ 1.6** — the recommendation's load-bearing part: five architectural arrangements, the count of governed conversions, the patient-safety reason for counting them, and the boundary of what counts as new infrastructure.
2. **§ 1.5** — why the choice of standard alone settles nothing: conformance is not interoperability, and profiling, compatibility and governance are three different questions.
3. **§ 1.1** — what is externally fixed by European and German law, and what it leaves open.

§ 1.3 shows where dependency arises and feeds the tender conditions; §§ 1.2 and 1.4 give each side its hearing and are frame rather than mechanism. The comparison of openEHR and HL7 FHIR as standards follows in § 2.

1.1 The Regulatory Framework in Europe

The EHDS Regulation (EU) 2025/327 has been in force since March 2025; technical specifications are being developed through Implementing Acts that are due by March 2027. MyHealth@EU currently relies on HL7 CDA and is moving towards HL7 FHIR specifications, including for the Patient Summary and e-Prescription/e-Dispense. In Germany, § 355 SGB V specifications use FHIR; Switzerland is pursuing the same direction through SwissHDS. No comparable EU-level regulatory role was identified for openEHR. This establishes regulatory relevance, not technical or clinical superiority.

How the regulation becomes a specification, and who writes it. Implementing Acts do not invent technical content, they reference specifications produced elsewhere. The instrument producing them is **EURIDICE**, European Interoperability Specifications for Digital Solutions in Healthcare, a joint initiative of HL7 Europe and IHE-Europe whose stated purpose is to create **implementable** and **testable** specifications. They are offered to the European Commission for reference in the EHDS Implementing Acts; none has been adopted yet (Appendix D-12).

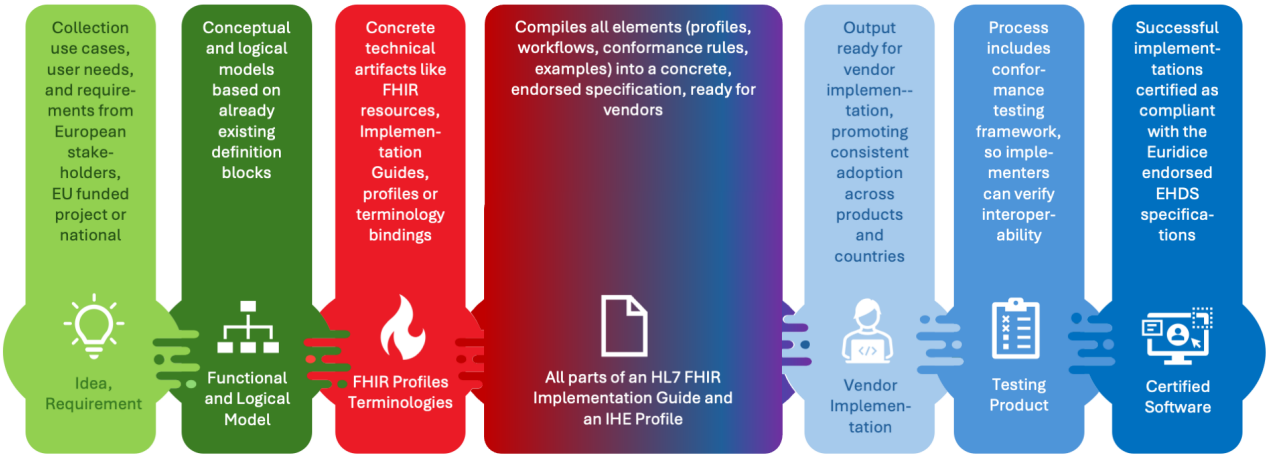

Figure 2 — Coverage of the Euridice initiative, from requirements through specifications to implementation and certified products.

Its artefacts cover the EU Patient Summary, the laboratory report, the (hospital) discharge report, the imaging report, medication prescription and dispense, manifest-based access to imaging objects and an EU Health Data API. The method combines HL7's FHIR expertise with the IHE profiling and testing methodology. Two of the guides were placed in coordinated ballot with a deadline of 30 April 2026 (Appendix D-12, where the author's own roles in one of the two bodies are disclosed).

What that settles and what it does not. The specifications currently being prepared for the Implementing Acts use FHIR implementation guides and the European profiling and testing machinery of HL7 Europe and IHE-Europe. This explains the exchange-boundary asymmetry in § 1.6.1 and makes a near-term departure from FHIR unlikely for the data exchanges concerned. It does not establish clinical quality or determine the internal persistence architecture. The European exchange direction is therefore substantially constrained, while the internal architecture remains open.

Against this stand real openEHR investments in several regions: Catalonia (CatSalut), Slovenia (Better platform), Norway (DIPS Arena), Sweden (Cambio), parts of the Finnish THL structures, and parts of the German Medical Informatics Initiative (HiGHmed) (Appendix D-3). What follows from this for EU standard-setting is a contested point and is recorded as one at § 1.7.

A word about NHS Wales. In June 2023 the Welsh Government mandated HL7 FHIR, via Welsh Health Circular, as the binding foundational standard for all NHS Wales bodies. The central national data platform (National Data Resource) stores clinical and administrative content in a **native, persistent FHIR store**, the Care Data Repository. Digital Health and Care Wales describes that store as the long-term structure of the data rather than a translation layer. With respect to the architectural question at issue here, Wales is the most prominent *national* FHIR-native CDR case in Europe (Appendix D-7.2).

An honest framing must acknowledge that neither standard covers the full breadth of interoperability. Blobel/Oemig (HL7-Mitteilungen 41/2018) extend the classic three-level HIMSS classification (technical, structural-syntactic, semantic) by four further levels: organisation and service, knowledge-based and skills-based interoperability.

Measured against that scale:

- **Levels 0 to 3, technical to semantic:** served by both FHIR and openEHR.
- **Level 4:** covered only in part.
- **Levels 5 and 6:** covered by neither standard.

The choice between FHIR and openEHR therefore plays out inside the same lower layers. Workflow orchestration, consent enforcement, cross-domain knowledge integration and the coordination of individual actors need complementary architectures and methods, whichever standard is chosen.

The same scale delimits what the European specification work described above can achieve. EURIDICE produces exchange specifications, and exchange specifications operate at the levels both standards already serve.

1.2 Arguments for FHIR That Deserve a Fair Hearing

Five strategic arguments support FHIR, and each carries weight regardless of an institution's legacy infrastructure. They are ordered by procurement logic: external obligations first, then the conditions of implementation, then the risks of dependency, and last the institutional setting, which is real but less decisive.

1. **Regulatory Connectivity**

   FHIR is the underlying standard for key EHDS-related data exchange initiatives, including MyHealth@EU and Germany's § 355 SGB V. Consequently, adopting an openEHR persistence layer requires ongoing cross-standard mapping into active FHIR exchange profiles: a transformation process that is not yet standardised across Europe.

   The anchoring has since been stated for the ordinance route as well. On 11 August 2026 the Interop Council asked that FHIR be the leading exchange format as a semantic and syntactic standard, with DICOM remaining admissible for imaging (Appendix E-11). That is a recommendation from the body constituted to advise on the ordinance and not a determination. It concerns exchange, and it names no persistence standard.

2. **Ecosystem, Tooling and Market Momentum**

   Independent of regulatory mandates, FHIR commands an expansive developer community and a mature, largely open-source ecosystem (including reference implementations, FSH/SUSHI profile authoring, instance validators, and regular connectathons). This lowers integration barriers, drives vendor competition, and reduces training overhead.

   Market momentum also develops beyond mandates. A survey of organisations planning SMART/HL7 Bulk Data interfaces found that only the EHR vendors themselves named regulation as their primary motivation. Payors, cloud providers and research bodies reported use cases of their own, for internal analytics, public health reporting and value-based care (Jones et al. 2021, self-reported plans, evidence grade C; Appendix G). Although based on a small 2020 US sample, it illustrates how regulatory mandates can seed an ecosystem that subsequently scales on market demand alone. It does not establish the same pattern for Europe.

3. **Workforce Availability and Delivery Risk**

   The available pool of FHIR developers, technical consultants and training programmes appears substantially broader than the corresponding pool for openEHR dual-level modelling. Skills in archetype development, template governance and clinical modelling remain comparatively specialised (*cf. Appendix C-4, FIKS recruiting*).

   The evidence supports scarcity in the documented cases and not a continent-wide quantitative comparison, because no comparative European labour-market count has been located. The advantage is also operational rather than technical. A broader workforce reduces recruitment dependency, delivery delay and reliance on individual specialists, and that becomes strategically relevant when capital investment reaches hundreds of millions of Euros. It does not establish that FHIR requires less clinical modelling expertise or produces better models.

4. **Vendor Diversity and Exit Risk**

   The operational openEHR repository market is served by a relatively small group of specialist vendors (the market is described at § 1.3). Claims that *"only one vendor will survive in practice"* are overdrawn. The smaller supplier and skills base nevertheless makes concentration, replacement and exit risk legitimate procurement concerns (*see § 1.3 for a detailed breakdown*).

   The relevant comparison is not the number of products carrying a standards label. What a procurer should compare is implementation substitutability, export mechanisms, model ownership, available skills and the cost of changing supplier.

5. **Standards-Development Process and Institutional Standing**

   HL7 International is an ANSI-accredited Standards Development Organisation (SDO) with a formal consensus and ballot process. The openEHR Foundation operates as an independent standards organisation without equivalent ANSI, ISO or DIN accreditation. The institutional basis of legitimacy is therefore asymmetric, which is not itself a measure of technical quality, openness or implementation fitness.

   openEHR's application to the Joint Initiative Council (JIC, Athens 2024) indicates a clear push to strengthen coordination with other standards organisations. It does not alter its accreditation status: the JIC is a coordinating assembly rather than an accrediting body, and several member organisations (such as IHE) are similarly unaccredited consortia.

Without adding a further argument, it should be noted here that historical feasibility is not automatically current performance. The hybrid architecture deployed by parts of Germany's Network of University Medicine (NUM), for instance, is frequently cited as a real-world success story. Available documentation covers only the 2020–2022 implementation period and does not establish its current operational state, comparative performance or long-term lifecycle cost (*see Appendix C-3*). The case belongs here because it limits a claim made for openEHR rather than because it supports FHIR: it should be cited as evidence of historical implementation feasibility, and not as evidence of present-day performance or economic superiority.

Addressing the Counter-Position: The Hybrid Ecosystem

A significant counter-argument to this assessment must be acknowledged. A 2026 systematic review in the *Journal of Biomedical Informatics* analysed 99 peer-reviewed studies (2021–2024) utilising FHIR, OMOP-CDM, or openEHR. It advocated for a **hybrid architectural model**: **openEHR for persistence, FHIR for data exchange, and OMOP-CDM for analytics** (Marfoglia et al. 2026, evidence grade B for its counts; Appendix E-9).

Although this assessment explicitly advises against multi-layered hybrid models for new greenfield infrastructure, Marfoglia et al. present a systematically assembled case from an independent research team.

Three critical caveats contextualise this counter-position without dismissing it:

- **Conceptual Design vs. Operational Evaluation:** The hybrid roles are derived from each standard's theoretical design intent rather than empirical evaluations of lifecycle costs or operational performance. These are gaps the review authors explicitly challenge future research to address.
- **Academic Focus vs. Production Systems:** 87% of the analysed studies originated in research environments. Commercial and vendor-maintained production deployments were excluded, which the authors themselves note: *"may underestimate the true level of real-world standard adoption."*
- **Unmeasured Data Transformation Loss:** The review highlights a key vulnerability in layered architectures: only 27 of the 99 studies evaluated how much source data was successfully retained during translation.

Consequently, the semantic loss incurred when moving data between openEHR, FHIR, and OMOP-CDM remains largely unquantified.

Conclusion: The hybrid counter-position stands as a credible design proposal, but it remains an unproven model awaiting empirical operational testing.

1.3 Vendor Lock-in — Four Dimensions That Must Be Treated Separately

The debate about vendor lock-in is conducted too narrowly in the materials. It focuses on vendor concentration in the openEHR CDR market and overlooks that lock-in arises in four structurally distinct dimensions, each with different consequences for procurement and standardisation decisions.

Figure 3 — Four dependencies, the layer each sits at and the instrument that removes it. Choosing a standard removes none of them by itself; the fourth has been largely absent from the debate (derivation: § 2.5.1).

Dimension 1: CDR Platform Lock-In

Market dynamics and technical mobility. openEHR CDRs bind customers to a small, specialised market (for example, Better, vitagroup (ehrbase), DIPS, Cambio, Code24), where a vendor failure carries immediate operational risks. FHIR CDRs bind to a broader, though similarly concentrated market (for example, Smile Digital Health, Health Samurai/Aidbox, Firely, hyperscaler PaaS).

Crucially, FHIR's promise of openness does not automatically resolve platform lock-in. While data migration between FHIR CDR products is technically feasible, it remains operationally demanding: implementation-specific extensions, search-parameter indexing, and subscription configurations are rarely portable.

A recent market inventory (Devitt, July 2026; Appendix A-1; evidence grade D) lists roughly thirty FHIR server products across hyperscalers, commercial CDRs, and open-source offerings. This highlights the key difference between the two ecosystems: the FHIR persistence market is markedly more populous than the openEHR CDR market.

As a result, lock-in on the FHIR side tends to be less acute, not least due to greater vendor competition and viable exit options, even if a few mature products dominate production use.

The "Red Hat problem": tooling-driven lock-in. Platform lock-in is typically viewed as a consequence of market concentration: few vendors create high vendor-survival risk. For openEHR, however, a second, structurally sharper dynamic applies.

The best-documented adoption hurdle for openEHR is the steep learning curve required for clinical modelling and its two-level architecture (§ 1.2, point 3; Appendix C-4, point 2; Appendix D-5). The most viable way to lower this barrier, and to make openEHR accessible to standard project teams, is for vendors to encapsulate this modelling complexity within simplified tools. Yet these user-friendly tools are almost invariably proprietary.

The underlying standard remains open in principle, but a single vendor owns the only practical implementation path. Solving the skill deficit thus risks inducing the exact vendor lock-in that openEHR was designed to prevent (Appendix E-6). While Red Hat successfully commercialised Linux without compromising open access, no equivalent model currently exists in the openEHR market.

Consequences for procurement and tenders. When an openEHR provider promises reduced complexity and lower skill requirements, procurers must scrutinise the trade-offs involved:

- **Artefact Portability:** Do the resulting templates, queries, and clinical data structures remain fully functional without the vendor's proprietary suite?
- **Data Sovereignty:** How is data ownership and extraction governed at contract termination?

This assessment is crucial: it determines whether the vendor neutrality that justified choosing openEHR in the first place is actually preserved in day-to-day operations (cf. Gartner caution, Appendix D-6).

Dimension 2: Hyperscaler Lock-In

FHIR-native cloud CDRs (Microsoft Azure Health Data Services, Google Cloud Healthcare API) shift the dependency to technology corporations with data interests of their own. Where EHDS or national rules require storage within the Union, this dimension becomes regulatorily relevant.

Dimension 3: Lock-In Through Proprietary HIS Internals

Proprietary internal architectures represent the dominant practical dependency in healthcare IT, even if current literature rarely quantifies its full scope.

Established Hospital Information System (HIS) products such as Epic, Oracle Health, Dedalus (ORBIS), CGM and Nexus continue to store clinical data in proprietary schemas tightly coupled to their application logic. Standardised interface layers (for example, FHIR APIs) do not alter these underlying persistence models, they merely expose a surface layer. Epic and Oracle Health, for instance, support extensive FHIR interfaces while keeping their core database architectures entirely proprietary.

Exceptions like DIPS Arena, built natively around openEHR, demonstrate that decoupled architectural designs are technically viable. They remain the exception rather than the market standard.

Implications for procurement. Because standard interfaces alone do not eliminate deep architectural lock-in, explicit contractual portability requirements (such as mandated extraction formats, schema transparency, and data exit rights) remain essential for health systems.

Strategic consequence for European hospitals. For most healthcare providers the immediate challenge is not replacing legacy systems but extracting structured, semantically enriched data from the existing HIS.

A pragmatic short-term strategy is a profiled FHIR exposure layer, often driven by mandates such as Germany's § 373 SGB V hospital-interface rules or the EHDS Regulation; it establishes a standardised exchange mechanism and leaves the proprietary core intact. As § 1.5 sets out, exposing a FHIR R4 API does not by itself guarantee interoperability, which depends on profiling depth, terminology bindings and data governance; the policy lever this implies is discussed at § 1.7.

Dimension 4: Lock-In via Model Governance and Local Terminologies

The lock-in consequence: from vendor dependency to model dependency. openEHR archetypes identify their nodes by internal codes (`at-codes`) whose meaning is defined in the archetype, not in the instance; the specification reading, with a worked example, is at § 2.5.1. When `at-codes` are used as values, the machine-processable meaning of persisted data can only be resolved via the corresponding archetype or Operational Template (OPT). Consequently, while openEHR data is vendor-neutral, **it is not automatically model-independent**.

In practical terms, the value proposition *"data outlives the application"* holds true only as long as the underlying model repository remains surviving, accessible, and version-controlled. Dependency shifts from the software vendor to model governance. It does not disappear. Clinical records are retained for decades, so an Operational Template used to commit a patient record in 2027 must remain continuously retrievable in 2057 to interpret that record correctly.

Consequences for tenders and procurement. Procurement specifications for openEHR implementations must explicitly govern model retention:

- **Ownership and Hosting:** Who maintains custody of the Operational Templates (OPT) across their lifecycle?
- **Version Control and Exportability:** In what formats are historical OPT versions exportable, and how long are they guaranteed to be preserved?
- **Post-Contract Availability:** How is access to the model repository maintained following contract expiration or vendor insolvency?

This requirement is absent from every recommendation reviewed here, including the Gartner Research Note (Appendix D-6) and the ÖGTelemed position paper (Appendix E-3). Both demand data sovereignty through vendor-neutral persistence while tacitly assuming permanent, unmanaged model availability.

1.4 Where openEHR May Offer a Specific Advantage

The points made for FHIR rest on breadth: regulation, ecosystem, workforce and supplier choice. The points for openEHR are more specific and, where they hold, deeper. They are ordered from the most specific structural advantage to the weakest supporting proposition.

1. **Shared Clinical Model Governance**

   The international openEHR Clinical Knowledge Manager (CKM) provides a formal review and publication process for the clinical archetypes submitted to it, progressing from *Draft* through *Team Review* to *Published*, with comments, revisions and reviewer participation recorded on the platform. This gives the community a visible mechanism for developing shared clinical models independently of individual applications and repository products.

   The international review of the *intraocular-pressure* archetype by the openEHR eye-care collaboration group illustrates the process. 32 specialists from 11 countries completed two formal review rounds on the CKM platform, and the comments are publicly auditable. The case documents the governance process, and not the archetype's clinical efficacy, adoption or interoperability in deployed systems (Fuhrmann et al., conference poster 2026; *evidence grade D*).

   FHIR can also be governed centrally and put through formal review, and the FHIR community itself warns that decentralised profiling can lead to "extensionitis" and localised model inconsistencies. openEHR ideally carries a shared clinical model library as a first-class part of its modelling approach, while FHIR commonly develops use-case-specific profiles in independently governed implementation guides. **The openEHR advantage arises only where projects actually reuse and maintain the shared archetypes rather than replacing them with local models.**

2. **Application-Independent Clinical Modelling**

   openEHR separates clinical content definitions from application and persistence implementation through archetypes and templates. This can support reuse of clinical models across applications and allow stored data to remain interpretable as applications change.

   The potential advantage is greater modelling continuity, and not automatic interoperability. It depends on disciplined template governance, terminology binding, version management and access to the complete model repository. Two systems using different templates, or different versions of the same template, are not interoperable merely because both conform to openEHR.

3. **Compatibility with FHIR-Based Exchange Architectures**

   The European Health Data Space (EHDS) framework does not exclude openEHR. An openEHR repository can operate behind a FHIR exchange layer, provided that the required transformations are implemented and clinically validated.

   Pedrera-Jiménez et al. (*JMIR 2023*; evidence grade C, authors include openEHR-side figures, Appendix G) conclude that openEHR, ISO 13606 and FHIR are functionally and technically complementary, with the optimal choice depending on system purpose. The Madrid INFOBANCO platform serves as a production example of this hybrid approach in practice.

   These sources establish architectural compatibility and implementation feasibility. They do not demonstrate that a hybrid architecture has lower lifecycle cost, better clinical outcomes or less governance effort than a FHIR-based alternative. **Compatibility is not an advantage over FHIR. It removes an objection, and it creates a mapping obligation.**

4. **Potential for Model and Data Portability**

   openEHR is designed to separate clinical models from repository implementations, and its specifications provide standardised artefacts for archetypes, templates, queries and record extraction. In principle this can make clinical content and its governing models less dependent on one application supplier.

   That potential should not be equated with achieved vendor neutrality. Product support for standard interfaces and export mechanisms varies, and operational portability also depends on modelling tools, intellectual-property rights, terminology licences, skills and contractual exit provisions. FHIR implementations create dependencies of their own through local profiles, extensions, terminology services and proprietary mapping layers.

   Lock-in is therefore possible in both ecosystems. The procurement question is whether data, models, mappings and terminology bindings can be exported, interpreted and maintained after a supplier change (*see § 1.3*).

5. **External Analyst Support Under Strict Preconditions**

   A Gartner Research Note (April 2017; *reviewed via summary — see Appendix D-6 for sourcing boundaries*) identifies openEHR as a viable architectural option for CIOs seeking to avoid single-vendor lock-in and to construct a long-term digital ecosystem. Gartner frames openEHR as a **complement to — not a replacement for — megasuite EHRs**, and makes the case conditional on preconditions:

   - A clear, long-term ecosystem vision
   - Sustained investment in clinical informatics skills
   - Mature governance

This is a historically relevant external assessment and not evidence of current market maturity, comparative effectiveness or successful implementation. It is nine years old, it was reviewed through a summary rather than the complete research publication, and that sourcing boundary remains explicit.

Institutional support no longer rests on Gartner alone. In 2026 a Delphi study commissioned by the Spanish Ministry of Health recommended openEHR as the most suitable standard for the national digital medical record, with FHIR at the exchange layer (Appendix E-12). That is the weightiest institutional endorsement the openEHR side holds to date, though with limits: it is an expert consensus that calls itself not evidence-based, its authorship overlaps with openEHR's leadership, openly declared, and it is a recommendation rather than an adopted strategy.

1.4.1 The Catalonia Case — Procurement and Outcomes Assessed Separately

Catalonia is one of the largest and most frequently cited regional openEHR deployments in Europe. Its procurement history is documented (chronology and source classes: Appendix D-2) and shows that an openEHR-based regional platform strategy is organisationally and legally feasible under European procurement frameworks; it matches the scenario Gartner (2017) names for national or regional programmes with a long-term horizon (Appendix D-6). What has not been published in independent, peer-reviewed form is any operational outcome: active institutions, record volumes, performance or quantified interoperability gains. The most-cited account is a Springer chapter whose lead author has since become CEO of openEHR International, which places it in institutional self-presentation within a peer-reviewed framework rather than in independent evidence (Appendix D-2).

Strategic Takeaway: as evidence of operational effectiveness, long-term cost-efficiency or technical superiority **the case is not currently robust**, and it should be cited with a qualifying caveat. *Catalonia proves that an openEHR path is politically and procurement-legally viable; whether it achieves its long-term clinical and economic objectives cannot yet be empirically judged.*

1.5 Conformance Is Not Interoperability

> Conformance and interoperability are separate properties: two systems can each pass a conformance test against the same base standard and still fail to work together. **The condition is that this holds only where profiles, versions, terminology bindings and process assumptions have themselves been harmonised and tested**, which is a governance task and not a property the standard supplies.

A distinction frequently blurred in policy debate is that **conformance** and **interoperability** are separate properties (Oemig et al., *HL7-Mitteilungen* 2018):

- **Conformance** means satisfying the syntax, constraints, and rules of a specific named specification, profile, and version.
- **Interoperability** means that two or more concrete systems can reliably exchange and correctly interpret information to execute a defined task.

Two systems may conform fully to the same base standard and still fail to interoperate if they implement different profile families, version variants, terminology bindings, or workflow assumptions.

Core Implications for Policy and Architecture

1. **"Mandating FHIR" is insufficient.** Selecting FHIR merely establishes a common technical baseline. Achieving true interoperability additionally requires harmonised implementation guides, globally or nationally resolvable terminologies, synchronised versioning, rigorous conformance testing, and cross-institutional governance.
2. **"openEHR is semantically richer, therefore more interoperable" is equally insufficient.** Greater modelling depth preserves more clinical context, but it contributes to interoperability **only** where the underlying archetypes, operational templates, terminology bindings, and interpretative logic are shared or mapped via

validated tooling. Semantic richness expands expressiveness; it does not inherently guarantee aligned interpretation.

Local Customisation: The Universality of the Problem

The widespread reliance on local extensions highlights this friction across all major standards. In a systematic review of published implementations (Marfoglia et al. 2026; Appendix E-9):

- **FHIR:** 23 out of 27 reported implementations required local adaptation rather than using the standard out-of-the-box.
- **OMOP-CDM:** 19 out of 44 implementations adapted the model, while 25 used it unchanged.
- **openEHR:** Split evenly (4 out of 8 studies required local adaptations).

The friction has since been quantified at the level of deployed systems. Across 68 US oncology sites, the probability that a common structured element arriving from another site was recognised at the receiving site averaged 0.68 between sites on the same vendor's product and 0.22 between vendors, although recommended standards existed for every element measured (Bernstam et al. 2022; Appendix D-11.3; evidence grade B for the scores). This shows that even "same standard, same vendor" buys only partial mutual intelligibility without governed mapping.

These figures reflect published research and demonstrate that **local customisation is predominant in reported FHIR deployments**. Extensibility enables adoption across heterogeneous clinical environments, but every ungoverned local extension creates a new compatibility boundary.

The OECD reaches a similar conclusion from a global policy perspective. Uneven standards adoption, vendor-specific API extensions, locked data models, fragmented governance, and a lack of rigorous conformance testing remain primary barriers to digital health transformation (OECD 2026; Appendix E-10). Simply selecting an open standard does not eliminate the need for centralised profile governance and binding conflict-resolution authorities.

Compatibility Conflicts Between Conformant Profiles: Two Documented Cases

Example A: Intra-National Profile Conflicts (Germany). The Core-Profile Working Group of the German Interop Council (March 2025) identified a direct contradiction within the national FHIR ecosystem (Appendix D-10.3):

- **Hospital Interface Profiles:** Mandate `Patient.gender` as a required element.
- **E-Prescription Profiles:** Prohibit `Patient.gender` on privacy grounds.

Both profile sets are fully conformant with FHIR, and both nominally rest on the same national base specification. This demonstrates that independently governed profile families create incompatible systems unless cross-profile conflicts are actively identified and resolved.

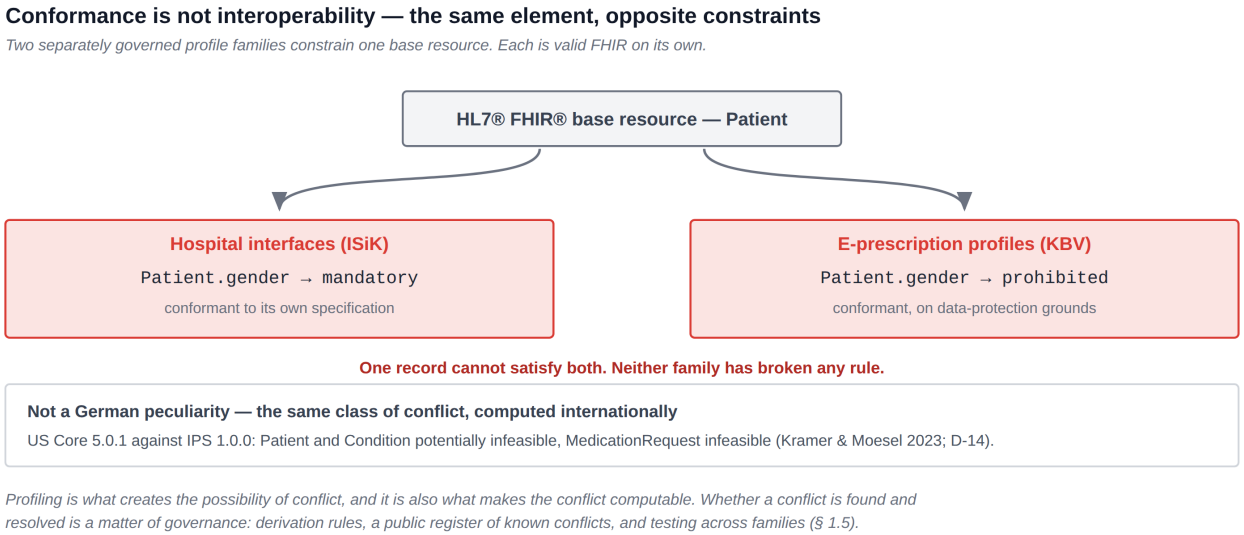

Figure 4 — The German case and its international twin in one picture: two separately governed families constrain the same base resource in opposite directions, each conformant to its own specification. Profiling makes the conflict possible and also makes it computable; governance decides whether it is found.

Example B: Cross-Border Profile Conflicts (US Core vs. IPS). Kramer & Moesel (2023) developed a predictive methodology to evaluate whether FHIR resources generated under one Implementation Guide (IG) can be parsed by a receiver enforcing another IG. They test **US Core 5.0.1** (sender) against the **International Patient Summary (IPS) 1.0.0** (receiver; Appendix D-14):

RESOURCE AREA	CONFORMANCE / MAPPING CONFLICT	INTEROPERABILITY VERDICT
`Patient`	IPS requires `birthDate` and mandates specific `contact.relationship` bindings; US Core does not enforce these requirements.	**Potentially Infeasible**
`Condition`	`clinicalStatus` is mandatory in IPS but optional in US Core, accompanied by divergent data types and terminology bindings.	**Potentially Infeasible**
`Medication`	US Core uses `MedicationRequest` with `reported=true` for patient-reported meds; IPS mandates the `MedicationStatement` resource for reported medications.	**Infeasible** (Base resource mismatch causes structural rejection)

Conclusion and Governance Implications

The examples above do not imply that FHIR or openEHR are flawed standards. Rather, they demonstrate that **conformance of sender and receiver to their respective implementation guides does not, on its own, establish operational interoperability.**

This distinction reframes the interpretation of profile collisions in two key ways:

a) Profile collisions are architectural in origin

The `Patient.gender` conflict is not a unique flaw of German governance. It is a fundamental property of **independently governed profiling**: whenever separate bodies constrain a shared base model for different operational purposes, the resulting profiles can be individually valid yet share no common conformant instance.

Testing each profile in isolation will never reveal this incompatibility — it requires explicit **cross-profile compatibility analysis**.

b) FHIR exposes collisions, governance must resolve them

FHIR is not unique in permitting profile conflicts. Its formal profiling framework makes constraints machine-readable. This allows both **collisions to be detected and automated comparisons** to be carried out.

What makes the German case notable is its institutional environment: divided specification responsibilities, conflicting statutory mandates, and historically missing cross-family reconciliation authorities. The Interop Council working group has documented the conflict and, with the creation of "base profiles", has already proposed solutions.

Procurement and Policy Takeaways

The possibility of incompatible specialisations cannot be eliminated without destroying the flexibility that makes open standards useful; it can be governed. For the bodies that determine specifications that means derivation from unified base profiles, formal rules for permissible constraints, automated cross-family comparison before release and a named authority for the conflicts that arise anyway (Document 6, items 13 and 14). For procurement it means accepting no isolated conformance certificate and requiring proof of cross-profile compatibility across every profile family the system must serve at once (§ 1.6.3).

The set logic that makes this computable, and where it is written down. The FHIR specification itself describes a profile in terms of the instances that satisfy it: each profile "*can be thought of in terms of the set of instances that conform to the profile*", and two profiles can stand in one of four relations to each other, non-overlapping, partly overlapping, one set contained in the other, or identical. It draws the practical consequence in the same place: an application creating content for two profiles must satisfy their intersection, while an application consuming content must accept their union [1]. Where the intersection is empty, as it is for `1..1` against `0..0`, no producer can serve both, and each side still passes only its own validator.

The capacity for conflicting constraints belongs to profiling. Detecting them is a compatibility question. Deciding whether they are prevented, accepted, mapped or resolved is a governance question.

In other words: governance has to require the comparison in order to obtain a basis for deciding what is to happen with the contradictions. Conformance testing verifies the resulting rules. Interoperability is demonstrated only where concrete systems can use the information for a defined purpose.

The openEHR symmetry

This governance issue applies equally to openEHR. Operational Templates (OPTs) maintained by different health authorities or vendor teams may constrain the underlying archetypes in mutually exclusive ways. A `Composition` valid under one OPT will not necessarily parse under another without transformation.

Currently, the published literature lacks a cross-template compatibility study for openEHR comparable to Kramer & Moesel's FHIR analysis. **This represents an evidence gap, not proof that openEHR inherently avoids template-level collisions.**

[1] HL7 FHIR R4 (v4.0.1), *Profiling*, § 5.1.0.22 *Supporting Multiple Profiles*. `hl7.org/fhir/R4/profiling.html`, retrieved 17 August 2026. The passage on incompatible rules at § 5.1.0.14 *Re-profiling and Re-slicing* concerns a derived profile and its own parent, not two sibling profiles of one base resource, and is deliberately not cited for the case discussed here.

For openEHR, the strongest available product-level evidence is used (Delussu et al. 2024; survey of 17 repositories / 11 vendor & developer responses; Appendix A-2):

- **Query Mechanisms:** 6 respondents declared support for Archetype Query Language (AQL), 4 reported SQL access, and 2 used proprietary formal query mechanisms (with overlapping support).
- **API & Serialisation:** 4 declared a mostly conformant openEHR REST API (1 under development). Only 1 declared direct native support for the standard EHR Extract (1 provided it via a separate product).
- **Interoperability Obstacles:** The survey revealed significant variance in Composition serialisation, supported Reference Model versions, and API implementations. The study authors concluded that these differences impede straightforward repository interoperability and highlighted the **absence of a common conformance-testing framework**.

Synthesis: Two Independent Empirical Findings

Read together, the German FHIR corpus analysis and the openEHR product survey validate that **conformance is distinct from interoperability**:

- **In German FHIR Specifications:** Independently governed profile families can impose incompatible constraints even though all are expressed within the same base standard (Appendix D-10.3), and the profile-corpus count shows the fragmented derivation structure that permits it (Document 5, D-10.7).
- **In the openEHR Market:** Commercial products marketed under the same specification family expose disparate combinations of query languages, APIs, serialisation formats, and extract mechanisms.

> **Core Strategic Takeaway**
>
> Adopting an open specification (FHIR or openEHR) establishes a common framework, but **not** operational interoperability. Achieving true interoperability requires:
>
> 1. Compatible and explicitly versioned constraints 2. Unified terminology bindings and shared models 3. Native, tested portability mechanisms 4. Cross-system testing and binding governance authorities

Implications for Architecture and Mappings. What these findings mean for the complexity, cost and validation of transformations across standards is explained in § 1.6.

1.6 Mapping — the Costs as the Core Economic Question, and the Patient-Safety Aspects

Every governed boundary between clinical models creates continuing specialist work: modelling on both sides, transformation logic, round-trip fidelity, version management, and clinical validation. **The number of boundaries is an architectural indicator and not a measure of cost or safety.** One badly governed transformation can be more hazardous than two well-controlled ones.

The economic case for the much-cited complementarity architecture, that is, an openEHR CDR (for storage) with FHIR interfaces (for exchange), depends on whether the benefits of its additional governed model boundary exceed its lifecycle cost. The available literature establishes that mappings require specialised and continuing work. It does not, so far, provide comparable ten-year cost measurements across architectures. The argument must therefore distinguish documented engineering tasks from unmeasured cost assumptions.

The Five Mapping Effort Vectors

The mapping effort in a dual-standard openEHR-to-FHIR architecture is concentrated across five operational areas:

- **Parallel Model Governance:** Clinical content must be represented in both an openEHR archetype/template and a corresponding FHIR profile. While existing models can be reused, maintaining two model families, and

governing the relationships between them, requires continuous cross-standard expertise (Pedrera-Jiménez et al. 2023).

- **Transformation Logic (ETL):** Mappings must reconcile paths, data types, cardinalities, terminology bindings, unit conversions, and contextual references. Simple fields map directly. Complex concepts, however, may require custom FHIR extensions, composite elements or calculated values, and where necessary explicit decisions to omit unmapped source content.
- **Semantic Round-Trip Fidelity:** Syntactically valid output does not guarantee that clinical intent is preserved. Mappings require clear rules for transformed versus omitted content. Bidirectional workflows additionally demand rigorous testing to ensure data can be transformed back and forth without clinically relevant loss.
- **Version and Dependency Management:** Archetypes, templates, FHIR releases, Implementation Guides, and terminologies evolve on independent schedules. Any update to underlying constraints or codes requires impact analysis, logic revisions, and revalidation across package dependencies (Document 5, D-10.7).
- **Testing and Clinical Validation:** Both source and target payloads require multi-layer verification: conformance checks against target specifications (such as EHDS implementation guides), terminology validation, expected-output testing, clinical reviews, and automated performance and round-trip testing.

This transformation burden is not unique to openEHR and FHIR. Mandl et al. (2020) evaluated the SMART/HL7 FHIR Bulk Data API. They note that traditional mappings from local health records into models like i2b2, PCORnet, and OMOP are inherently expensive, slow, and reliant on continuous custom work by specialised teams.

Their study lacks a comparative cost analysis, but it reinforces a fundamental principle: **every governed boundary between source and target models incurs ongoing specialist effort**. The strategic question is not whether these boundaries consume resources, but which ones deliver sufficient value to justify their lifecycle cost.

Similarly, the OECD frames data mapping as a recurring overhead throughout the health data lifecycle, warning that model conversions risk compromising semantic context, granularity, and fitness for purpose (OECD 2026; Appendix E-10).

The report also claims that mapping inherently degrades data quality, but its evidence barely supports that. A well-governed transformation need not compromise clinical intent. It nevertheless highlights a critical architectural reality: **every schema conversion represents a point where semantic integrity must be explicitly specified, tested, and governed**. Notably, the report does not quantify the exact financial cost or data loss associated with any specific FHIR, openEHR, or OMOP implementation path.

Secondary Use and OMOP Projection

Analytics and population health introduce another distinct target model: the **OMOP Common Data Model (CDM)**. Regardless of whether operational data resides in a proprietary HIS, a FHIR server, or an openEHR CDR, cohort analysis requires projecting data into OMOP's fixed relational schema.

Multi-stage pipelines exist, such as mapping openEHR microbiology findings through FHIR into OMOP (Rinaldi and Thun 2021). But FHIR is not a mandatory intermediate step: direct openEHR-to-OMOP and FHIR-to-OMOP pipelines are equally viable. Consequently, economic evaluations must compare **complete end-to-end architectures**, rather than counting isolated transformation steps in a single pipeline.

Field experience and evidence limits. Real-world deployments corroborate the operational weight of ETL maintenance:

- **Anecdotes and Project Data:** High ETL effort and mapping errors reported in Germany's NUM initiative align with Scandinavian project reports.

- **Socio-Technical Friction:** The FIKS longitudinal study (Christensen and Ellingsen, 2016; Appendix D-5) documents widespread socio-technical challenges in large-scale model-driven development.
- **Integration Contingencies:** The Swiss EY study identifies legacy integration as a primary execution risk. However, precise empirical data regarding long-term maintenance costs, defect rates, and upgrade overhead remains unpublished.

Strategic Conclusion and Procurement Test

Mapping is not a one-off project. It is a durable architectural decision.

An openEHR CDR paired with a FHIR exchange layer introduces a durable, dual-standard boundary that requires continuous lifecycle maintenance. Conversely, a FHIR-native CDR eliminates this specific cross-standard boundary if its core data natively aligns with exchange profiles, though it does not eliminate the mapping that takes data in from proprietary HIS (ingest mapping), profile reconciliation, or OMOP analytics transformations.

The Strategic Decision Framework

A complete 10-year Total Cost of Ownership (TCO) evaluation must weigh modelling, ETL authoring, terminology maintenance, clinical validation, version upgrades, tooling, and exit costs against the architectural benefits of an extra model layer.

Because published longitudinal TCO data does not yet exist, healthcare organisations face a straightforward **procurement test**:

> **The Procurement Rule:** Introducing a dual openEHR-plus-FHIR model boundary is justified **only** when the specified clinical modelling depth or governance benefits deliver value that explicitly outweighs its ongoing lifecycle maintenance cost.

1.6.1 Operational Architecture Types and Boundary Methodology

The core operational decision facing health systems is not a simple choice between FHIR and openEHR. Rather, it concerns where clinical content is committed authoritatively, which governed model serves as the primary system of record, and which representation is enforced at the exchange boundary.

Evaluating these criteria yields **five distinct operational architecture types** (ideal types rather than an exhaustive taxonomy of all clinical deployments).

Conversion Counting Methodology

To evaluate structural complexity across these architectures, this assessment tracks **conversions between independently specified, versioned, and governed clinical models**.

- **Exclusions:** Transient software layers, user-interface view models, temporary caches, and derived database indices that can be dynamically reconstructed from the authoritative record are **not** counted as model boundaries.
- **Inclusions:** Any transformation between two distinct governed specifications where semantic alignment must be established, clinically validated, and maintained across schema version updates.

Conversion Classifications

1. **Transient Conversions:** The target model payload is dynamically generated for a specific transactional exchange or query and discarded after transmission.
2. **Persisted Conversions:** The transformed content is stored and maintained long-term in a secondary governed schema or repository.

Metric Limitation: The resulting conversion count serves as a **structural orientation**, not a complete Total Cost of Ownership (TCO) calculation. In practice, mapping breadth, semantic friction, validation requirements, artefact reuse rates, and schema update cadences often exert a greater influence on lifecycle costs than the raw count of model boundaries alone.

Please note, that the comparison assumes FHIR as the mandatory exchange representation for external communication. Reported figures reflect minimum cross-model conversions; additional intra-standard conversions may occur when reconciling incompatible FHIR versions or divergent Implementation Guides.

	AUTHORITATIVE CAPTURE AND PERSISTENCE	PATH TO THE REQUIRED FHIR REPRESENTATION	MINIMUM GOVERNED CONVERSIONS	EVIDENCE STATUS
A	Clinical content is captured and retained in a proprietary model	Proprietary model → FHIR when required	**1, transient**	Common architecture in existing installations
B	Clinical content remains authoritative in a proprietary source and is replicated into a persistent FHIR repository	Proprietary model → FHIR store	**1, persisted**	Demonstrated at infrastructural scale, generally with data populated from source systems
C	Clinical content is captured and retained authoritatively in FHIR	Direct use of the authoritative FHIR representation	**0**, if compatible with the required exchange profile	Testable requirement; one vendor declares it for its product, but no independently verified instance was identified (Document 4, D-7.8)
D	Clinical content originates in a proprietary system and is replicated into openEHR before FHIR is produced	Proprietary model → openEHR → FHIR	**2:** one persisted and one transient	Documented in particular integration architectures, not as a universal openEHR deployment pattern
E	Clinical content is captured and retained authoritatively in openEHR	openEHR → FHIR when required	**1, transient**	A major Norwegian development programme is documented; broad availability outside the Norwegian market is not established

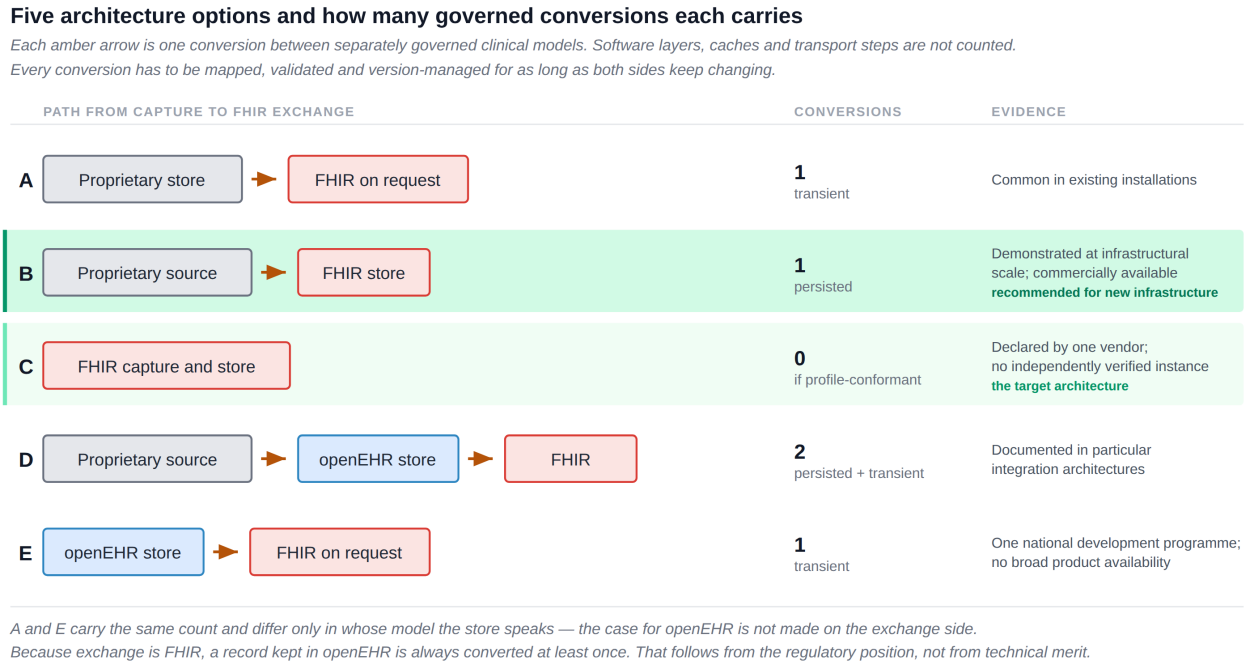

Figure 5 — Where the clinical record is written, where it is kept, and the minimum number of governed conversions needed for the required exchange representation. The count indicates lifecycle obligations and potential safety-relevant failure points; it is not a measure of cost or safety, since one poorly controlled conversion can be more dangerous than two well-controlled ones (§§ 1.6 — 1.6.2).

Equivalence Fallacy (Option A vs. Option E)

Options A and E share the same minimum count of cross-standard conversions (**1, transient**) when external FHIR export is required. To an external receiving system, a conformant FHIR payload reveals nothing about whether its authoritative source was a legacy proprietary database or an openEHR CDR. However, this does not make the architectures equivalent: their underlying data models, query capabilities (for example, AQL vs. SQL/FHIR search), vendor lock-in risks, and long-term portability conditions differ radically.

The imbalance in data conversion (the conversion floor)

Where the legislator mandates FHIR for data exchange, there is a built-in disadvantage in data conversion: the conversion floor is asymmetric:

- **The data are stored in openEHR (Options D and E):** There is a hard floor: the data must **always be converted into FHIR at least once**. No optimisation can avoid that effort.
- **The data are persisted FHIR-natively (Option C):** Here the theoretical floor of **0 conversions** applies. However, this zero-conversion state holds **only** if the stored FHIR resource is natively compatible with the destination exchange profile. Incompatible FHIR versions or divergent national Implementation Guides reintroduce a governed transformation layer.

For new infrastructure, the principal comparison is therefore usually between Options B and D. Option B is a persistent FHIR representation reached through one conversion. Option D is a persistent openEHR representation followed by a second conversion to FHIR, justified if openEHR modelling depth and governance value outweigh two transformation layers.

For existing installations, Option A often remains the starting point. Replicating legacy source data into a persistent FHIR store as in B is also an option.

Option C is the structural minimum, and it is claimed in practice. nursIT Institute states that its careIT products write clinical content directly as FHIR resources through the FHIR client with nothing in between, that the repository is interchangeable wherever it supports FHIR R4, and that eight hospitals run on a third-party repository on that basis (Appendix D-7.8). **The statement about its own architecture is the vendor's own and therefore carries evidence grade D.** The vendor's account can, however, be regarded as a demanding solution for the complete FHIR record, without that having been independently validated. More on this in the following section.

Procurement Criteria for Claims of FHIR-Native Capture (Option C)

If a vendor claims to deliver an Option C architecture (direct, authoritative FHIR capture of the core clinical data without internal proprietary schemas), tenders must mandate proof of the following:

1. **Native Validation:** Authoritative clinical content is committed and validated directly against specified target FHIR profiles at runtime.
2. **Schema Integrity:** Internal proprietary structures store only reconstructible application data such as UI, workflow or presentation state, not core clinical data.
3. **Enterprise Capabilities:** The native FHIR repository itself provides version management, provenance, audit logging and transactional integrity (ACID, glossary F-4). A mere exchange endpoint in front of someone else's data store does not.
4. **Direct Export:** Exchange profiles are generated directly from authoritative holdings without an undisclosed clinical translation layer.
5. **Demonstrated Reconstruction:** Full clinical record reconstruction and profile conformance are successfully verified during acceptance testing.

A selection of practical cases and their attribution

- **HiGHmed (Option D for documented paths only):** an openEHR clinical data repository placed ahead of the FHIR-based core dataset of the Medical Informatics Initiative (Ammon et al. 2024; Appendix D-7.4). The publications document the FHIR process and sharing layer, and a site-level openEHR-to-FHIR mapping is published for one domain only (Rinaldi & Thun 2021; Appendix D-8.3), so HiGHmed is recorded as Option D for the documented paths and not as a consortium-wide architecture.
- **DIPS Arena (Option E):** native openEHR development inside a major hospital information system programme. The Norwegian deployment shows technical viability; its market figures reflect vendor presence in Norway, not adoption across broader European markets (Appendix D-5).

Strategic Summary

Conversion counts narrow the architectural decision but do not dictate it:

- **Option D** is justified when the long-term clinical modelling depth, archetype reuse, and vendor-neutral governance of openEHR explicitly exceed the cost and maintenance of an extra mapping layer.
- **Option B** provides a structurally simpler, lower-overhead architecture whenever a persistent FHIR data model is sufficient for internal clinical and analytical requirements.
- **Option C** offers the fewest conversions under a FHIR exchange mandate, but its operational feasibility must be rigorously demonstrated in acceptance testing rather than assumed from marketing materials.

1.6.2 Patient Safety as a Distinct Dimension of the Mapping Question

Mapping is commonly treated as a cost issue in procurement. It is also a patient-safety issue. A transformation between clinical models, terminologies or systems can alter, omit or misrepresent information on which clinical decisions depend. The available evidence establishes this possibility and the resulting need for clinical validation. It does not quantify the patient harm attributable to particular architectures.

What the evidence establishes

The classes below are this assessment's own grouping of mapping deficits and defects against the background of patient safety. The studies named next (D'Amore, Wright, see Appendix D-11) evidence the first and the last; the others are stated as categories rather than as findings:

- **Semantic Heterogeneity and Errors:** Exchanged clinical documents often carry conflicting or misaligned semantic structures despite passing basic validation.
- **Terminology Granularity Loss:** In mapping between disparate terminology systems (or local codes), essential clinical specificity can be lost (for example, mapping a detailed SNOMED CT diagnosis to a broader ICD-10 category).
- **Omissions During ETL Execution:** Complex transformations may silently drop unmapped attributes, state flags, or protocol contexts that lack direct target elements.
- **Presentation and Routing Defects:** Interface layer errors can misalign, obscure, or delay vital clinical details before they reach clinician dashboards.
- **System Migration Failures:** Replacing underlying storage models or migrating historical data can alter contextual data dependencies, altering safety-critical application behaviours.

These classes should not be conflated. An interface incident is not necessarily a mapping error, and an interoperability defect does not necessarily reach a patient. Together, however, they show why technical conformance testing alone is insufficient for safety-critical information.

The study by D'Amore et al. (2014; Appendix D-11.1; evidence grade B) illustrates the range of semantic defects in standardised exchange payloads:

- **Sample & Methodology:** The authors evaluated 91 C-CDA documents across 21 distinct health-IT systems, conducting detailed analysis on 21 representative samples (1 document per vendor/system).
- **Observed Defects:** Across the 21 analysed samples, the study recorded **615 total instances of error and permissible variation**, including **142 cases categorised specifically as terminology misuse or omission**.

Key Takeaway: The findings demonstrate widespread semantic instability in multi-vendor exchange payloads. The documents were test-environment exports from self-selected certified technologies, and only the first instance of each problem type per document was counted, so the figure is a count of error types and not a rate. The study does not trace those observations directly to adverse patient outcomes or quantify harm attributable to specific pipeline steps.

What follows for architectural comparison

Each governed model conversion introduces another mapping that must be specified, clinically validated, versioned and monitored. Additional conversions therefore increase the number of potential failure points and the validation burden.

The evidence assembled here does **not** support the proposition that an architecture with two conversions produces measurably more patient harm in practice than one with a single conversion. The number of conversions is not, by itself, a measure of residual patient risk. Mappings differ in scope, complexity, transparency, tooling and reuse. A well-governed two-stage path may perform better than a single opaque transformation. The defensible architectural conclusion is therefore limited:

An additional transformation boundary requires an additional safety justification and validation regime. It should not be introduced without a demonstrated clinical or operational benefit.

This principle is standard-neutral. It applies to proprietary-to-FHIR, proprietary-to-openEHR, openEHR-to-FHIR and transformations between incompatible FHIR versions or profiles. Semantic depth in the source model may improve the information available to a mapping, but it does not remove the obligation to validate the target representation.

Data portability is not functional portability

A distinct patient-safety risk emerges when one clinical system replaces another. Even when a patient record is migrated with complete structural accuracy, safety-critical system behaviours can alter dramatically.

The Risk of Behavioral Drift: Evidence from System Migration. Wright et al. (2018; Appendix D-11.2; evidence grade A for what it measured at the site it measured) evaluated the replacement of a custom, locally developed electronic health record with a major commercial system at Brigham and Women's Hospital. Following the migration across 3,277 outpatient prescribers:

- **Alert Volume Escalation:** Interruptive drug — drug interaction (DDI) alerts rose from about 2 to about 11 per 100 orders, roughly **sixfold**.
- **Acceptance Rate Collapse:** Overall alert acceptance fell sharply. For the most severe alert tier, user acceptance dropped from **100% to 8.4%**.

Please note, that the pre-migration 100% acceptance rate reflected an un-overrideable "hard stop" that physically blocked orders, whereas the target system implemented a discretionary override workflow.

The study did not measure adverse drug events, compare specific standard architectures such as FHIR against openEHR, or allow direct quantitative extrapolation to other sites. Its core insight nevertheless remains vital for procurement and architecture planning: **preserving clinical data does not guarantee the preservation of clinical safety behaviours.**

Procurement strategies, tender specifications, and migration plans must explicitly decouple data structures from functional behaviours across three operational phases:

> **A migrated record is not a migrated safety net.** Alert severity, timing, hard stops, override rules, acknowledgement requirements and workflow integration may be implemented in application configuration, decision-support services or proprietary logic. Standards may represent parts of this logic, but neither a standardised repository nor a data export guarantees that the behaviour will be reproduced in another system.
>
> **Distinguish three things in the migration plan:** data migration, including structure, terminology, provenance and context; functional migration, including alerts, order controls, workflow constraints and other safety mechanisms; operational validation, establishing that the target system behaves safely in its intended clinical setting.
>
> **Count what has to survive the move, and do it before the contract is signed.** Specify the intended equivalents in the target system, test them, have them clinically accepted, and name explicitly before go-live every function that cannot be reproduced.

Testable Procurement Requirements for Data Mappings and Migrations

A generic tender requirement mandating *"100% mapping accuracy"* is operationalised poorly unless its underlying denominator, test domain, and definition of accuracy are explicitly specified. While universal accuracy cannot be demonstrated across an open-ended clinical domain, **complete accuracy can — and should — be required for a bounded, clinically approved acceptance set.**

For safety-critical domains such as medications, allergies and the like, procurement specifications must mandate:

- **Version-Controlled Acceptance Test Sets:** A clinically validated test corpus covering the specific concepts, complex combinations, and operational edge cases expected in production.
- **Explicit Semantic Preservation Criteria:** Mandatory retention of clinical intent across critical dimensions such as identity, status, certainty, temporality, dosage, route, units, negation, provenance, and contextual qualifiers.
- **No Unresolved Semantic Loss:** No unresolved loss of safety-critical meaning across the specified acceptance set prior to deployment.
- **Governed Edge-Case Handling:** Documented handling rules for unmappable or ambiguous content, defining both its user-facing presentation and its downstream impact on clinical decision support (CDS) rules.
- **Prohibition of Silent Omissions:** A strict structural ban on the silent dropping or unnotified truncation of defined safety-critical data elements.
- **Multi-Layered Validation:** Formal clinical review and semantic validation alongside standard syntactic and schema conformance testing.
- **Continuous Regression Testing:** Automated revalidation triggered whenever a source model, target profile, terminology set, mapping table, or application logic updates.
- **End-to-End Workflow Acceptance:** Testing executed directly within the receiving user interface and clinical workflow, rather than validating the isolated exchange payload alone.
- **Post-Implementation Surveillance:** Dedicated post-go-live monitoring and incident management protocols capable of isolating, reporting, and investigating mapping- and migration-related defects.

"No unresolved loss" and a numerical acceptance threshold are complementary. The first defines the required clinical outcome. The second makes performance against the agreed test set measurable. Neither is credible without a representative test corpus and explicit clinical acceptance criteria.

Conclusion

> Data mappings and system migrations must be governed as **patient-safety controls**, rather than treated as simple technical deliverables or transient integration costs.

Every additional model boundary introduces a transformation surface that demands its own dedicated validation, maintenance, and regression testing. While an architecture featuring extra model boundaries does not automatically cause patient harm, **the clinical and governance benefits of each additional transformation layer must explicitly justify its ongoing validation burden and residual operational risk.**

1.6.3 What Counts as New Infrastructure — the Decision Boundary

The recommendation carried in Part I and in the table of § 1.6.1 applies to *new infrastructure*. What does that mean in concrete terms? Four cases begin from different transformation baselines: a greenfield regional repository, a replacement hospital information system, a secondary-use platform, and an additional clinical data repository placed beside a running system. The layer-counting argument of § 1.6.1 and § 1.6.2 does not give the same answer for all four. This section states the boundary, and states where the recommendation stops.

The definition used here. New infrastructure means a data-persistence layer being procured or built for which no committed persistence model is yet in production for the data in question. The test attaches to the persistence layer and is not a budgetary question. A project that is new in procurement terms but writes into an established repository is not new infrastructure in this sense. A replacement is new infrastructure to the extent that its persistence model is genuinely open at the point of award.

Four starting positions.

STARTING POSITION	TRANSFORMATION BASELINE BEFORE ANY CHOICE IS MADE	HOW FAR THE RECOMMENDATION APPLIES
Greenfield primary capture — a regional or national repository with no incumbent persistence model	None. The persistence choice determines the whole downstream chain.	In full. This is the case the recommendation is about, and the case in which the layer-counting argument has its greatest force.
Replacement of a hospital information system or of an existing repository	One migration of the historical data, unavoidable and non-recurring, plus whatever the replacement product's own model imposes.	In full where the persistence model is genuinely open at award. The one-off migration is a project cost and not a standing transformation layer; counting it as a recurring layer overstates the case against either option.
Secondary-use platform fed from primary systems	At least one transformation whatever is chosen, because the data are replicated rather than captured at source.	Weakly. See the limit stated below: where the data arrive already transformed, layer count discriminates less between the options, and modelling depth and query requirements carry correspondingly more weight.
Additional repository beside a running primary system	At least one transformation, because the primary system's model is fixed and cannot be renegotiated by this procurement.	Partly. The question is no longer which model avoids the first layer, but which target model minimises the layers that follow it, given the external interfaces that are mandatory.

The criteria, and what each is evidence of. Two questions are already settled by the table above: whether the persistence choice is open at all, and whether the data are captured at source or arrive replicated. Where the decision is open, five criteria decide which way it points. Each is stated together with what it actually bears on. That is necessary because several of them are routinely used in tender documents to support a conclusion they do not support.

CRITERION	WHAT IT BEARS ON	WHICH WAY IT POINTS	CAUTION
Existing openEHR investment	Migration cost and available skills, not architectural fitness	Toward retaining the incumbent model where the investment is substantial and in production	A sunk investment is a budget argument. Presenting it as an architecture argument is a category error — see below
Mandatory external FHIR interfaces	The number of layers downstream of persistence	Toward FHIR-native persistence where such interfaces are extensive and binding	Where the external obligation is narrow and stable, the exposure layer is a bounded cost rather than a structural one
Required modelling depth	Whether the domain requirements exceed what profiled FHIR	Toward openEHR or ISO 13606 where clinical modelling depth and	The requirement has to be demonstrated against the domain, not asserted. § 1.6 records that round-trip transformation loses

CRITERION	WHAT IT BEARS ON	WHICH WAY IT POINTS	CAUTION
	expresses without extension overload	archetype governance are genuinely required	context unless profiles are overloaded; that argument runs in both directions
Number of unavoidable upstream transformations	The residual value of avoiding one further transformation	Toward the model that adds fewest further layers	This is the criterion the document gives most weight to where the persistence choice is open, and § 1.6.2 gives the patient-safety reason. It discriminates only where the count actually differs between the options
Evidence required for clinical safety and total cost	What a tender must ask for, whichever way it decides	Toward the option whose supplier can meet the safety conditions of § 1.6.2 and produce a ten-year cost series	Neither option is evidentially established for outcomes (Part I, Finding 3). A tender that requires outcome evidence for one standard and not the other is not neutral

Existing investment is a budget argument, not an architecture argument. An installed openEHR base with archetype governance, trained staff and running clinical content is a substantial asset, and writing it off is a real cost that a procurement may legitimately decline to incur. That is a decision about money and capability. It is not evidence that the architecture fits better, and the two are regularly conflated in practice.

The same discipline applies here as in Part I, Finding 3. A decision may rest on a ground of this kind honestly, provided it is stated as the ground it is.

Where the data are replicated, the core argument of this document discriminates less. The layer-counting argument assumes that the persistence choice determines whether the first transformation happens at all. In a secondary-use platform fed by extraction from primary systems, and in a repository placed beside a running system, that assumption does not hold: one transformation exists before the choice is made. What remains is the difference in the layers that follow. That difference is smaller and more dependent on the specific interface obligations.

A statutory exchange obligation is not a statutory persistence obligation. Of these criteria, *mandatory external FHIR interfaces* is the one most often carried further in tender reasoning than it will bear. The current German instance is a draft.

The Act on Data and Digital Innovation in Health Care was approved by the Federal Cabinet on 15 July 2026. Its § 386a SGB V obliges manufacturers of information-technology systems to make health data available without delay and free of charge in an interoperable format, and to enable access to it. The duty arises on request of those providing care who use the system, and it concerns that provider's patients' data. The format is fixed by ordinance rather than named in the statute.

The duty is addressed to manufacturers, and it is a duty to release and to grant access. § 386c prohibits a manufacturer from releasing data in a format that prevents complete semantic reconstruction by another admitted system, and from withholding data elements that are not among the bindingly determined requirements but are needed for the documentation and archiving duties that fall on those providing care. That second prohibition reaches past the prescribed exchange format, and it is model-neutral: it binds an openEHR-based system exactly as it binds a FHIR-based one.

The statutory advisory body has since drawn a third line. On 11 August 2026 the Interop Council, constituted under § 385 SGB V, asked that a statute prescribe semantically standardised data holding at the *logical* layer, so that data are held with standardised semantic annotation, while expressly asking that no statutory requirements be placed on databases or system architectures (Appendix E-11). On that reading a statute would prescribe that a blood pressure is annotated with a published concept and a defined unit, and would prescribe nothing about the store beneath it.

The duty asked for is narrower than its opening sentence. The Council would bind manufacturers only where a data content has been determined bindingly, naming units of measurement and observations as its examples, and it proposes that the determined set start as a first mandatory core derived from the European Patient Summary and be extended in stages.

What follows for a tender, and what does not. The Council advises; the competence centre proposes; a binding format is set by ordinance. A procurement today is bound by the operative text and not by the position: an exchange obligation constrains the exposure layer and does not fix the persistence model. What the position does change is the direction a procurer should expect the requirements to move in. If the logical-semantic duty arrives in the ordinance, it will be satisfied by annotation discipline and resolvable terminology bindings, which § 2.5.2 already treats as a testable requirement, and **not** by the choice of a persistence standard. The paper adds that the associated data models should in future be standardised by statute, naming no instrument, body or model; that is a proposal and not a plan (Appendix E-11). The Council's paper names FHIR, and names it for the exchange format alone. It names no persistence standard at all, openEHR included.

The factual record is at Appendix D-10.3.

Where new infrastructure is being built rather than connected, the decisive question is not which store is bought but **which structures hold the clinical content authoritatively**. A vendor keeping the clinical record in FHIR and using proprietary structures for application state, workflow, user-interface concerns and derived indices is at Option C. A vendor keeping a proprietary clinical model and generating FHIR from it is at Option A or B, however complete the FHIR output.

> **Tender condition — the authoritative-record test.**
>
> **Requirement:** the clinical system of record is FHIR, and every proprietary structure that holds clinical content is derived from that record and reconstructible from it. Application state, workflow, user-interface concerns and derived indices lie expressly outside the requirement.
>
> **Evidence with the bid:** a statement of which structures hold clinical content authoritatively, the derivation of each proprietary structure holding clinical content from the FHIR resources it rests on, and the procedure by which those structures are rebuilt.
>
> **Test at acceptance:** rebuild them from the FHIR record alone against a representative clinical data set, and show that the system behaves identically on that data.
>
> **Limits:** the condition applies where new infrastructure is built rather than connected (§ 1.6.3), and no comprehensive product is documented as meeting it today (§ 1.6.1; Appendix D-7.7 to D-7.9).

1.7 Open Strategic Points of Contention

1. The Binding Nature of European Standard-Setting. A hard EU mandate enforcing "FHIR only" across both exchange and persistence would be regulatorily clean, but it risks devaluing substantial regional and institutional openEHR investments. Conversely, open-model frameworks (such as SwissHDS) shift the harmonisation burden directly onto national and regional health authorities, creating local integration overhead.

2. Political Economy of the Complementarity Narrative. The political narrative framing openEHR and FHIR as complementary tools obscures a fundamental market tension:

- **FHIR Stakeholders** view the "openEHR for persistence, FHIR for exchange" argument as a rhetorical Trojan horse designed to mandate parallel openEHR repositories alongside existing databases.
- **openEHR Stakeholders** fear that defaulting to FHIR-native storage will erode domain-driven clinical modelling competence and reduce records to shallow exchange profiles.

Both concerns are voiced in the position materials examined in Appendix E; neither is measured.

3. Economics of Mapping Overhead. As long as the lifecycle costs of the cross-standard mapping layer remain hidden within general integration budgets, the complementarity narrative easily prevails in policy debates. **Mandatory 10-Year Total Cost of Ownership (TCO) comparisons** — factoring in dual-model governance, continuous ETL maintenance and clinical revalidation, would fundamentally alter this balance.

FHIR Exposure as a Policy Lever: Potential and Limits. Mandating either openEHR or FHIR-native storage as the internal persistence core for legacy Hospital Information Systems (HIS) is politically and economically unfeasible in the short term. The most pragmatic regulatory lever is therefore **obliging every HIS vendor to expose a profiled FHIR server interface**.

The Need for Binding Conformance Regimes. For an interface mandate to be effective, it requires a rigorous conformance-testing regime, a certification infrastructure and centralised governance: machinery that remains fragmented across most of Europe.

Germany demonstrates that such an infrastructure is buildable. Under **§§ 385 and 387 SGB V**, gematik's Competence Centre for Interoperability operates a mandatory conformity assessment with published test suites, issued confirmations and a public positive list. As of 11 December 2025 that list recorded 141 assessed systems for the ePA medication service (self-reported by the operating body, evidence grade D; Document 5, D-10.8), with ISiK Stage 5 announced as a next subject.

However, Germany's experience also highlights the limitations of this approach: **passing a test suite is a statement about the test, not about general clinical usability or cross-system interoperability (§ 1.5)**. Testing remains confined to defined artefact sets rather than guaranteeing broad semantic compatibility.

Conclusion: Market Economics and the Persistence of Mapping

The valid question is why, after four decades of HL7 specifications, healthcare architectures still rely heavily on complex mapping layers. The answer lies not in technical limitations in the standards but in the **economic logic of the HIS market**:

> **Core Regulatory Insight:** Data transformation boundaries persist because no single market participant bears their total lifecycle cost alone. The financial and clinical risks of mapping are fragmented across health authorities, IT departments, and clinicians. In this assessment's judgement, that systemic market failure is why data architecture and mapping governance must be addressed through **binding national regulation**, rather than delegated to isolated institutional procurements.

2. Technical and Professional Level

What this part establishes, and what it assumes. Part I states the findings and the recommendation, and § 1 derives them from the regulatory and economic situation. This second part examines the two standards themselves: what each was built for, how each binds terminology, what each governs, and where the debate compares them on grounds that do not hold. It can be read on its own, with one qualification.

Three things are established earlier and are used here without being restated. The first is the five architectural arrangements, together with the count of independently governed models the content must cross (§ 1.6.1). The second is the distinction between conformance and interoperability (§ 1.5). The third is the rule that an evidence grade attaches to a proposition together with its source rather than to a publisher (Document 4, Appendix G). Abbreviations are resolved in the glossary at Appendix F.

Here too the order is the order of the argument, not the order of weight. §§ 2.1 to 2.3 set the frame, the mechanism sits in the middle, §§ 2.6 and 2.7 clear away objections and add the research perspective. Read in this order if time is short:

1. **§ 2.5.2** — binding strength, where the difference between the standards actually sits, with the corpus finding on the safety-critical subset; the sharpest technical finding of this document.
2. **§ 2.5.1** — the self-description of instance data and openEHR's model dependency, which § 1.3 Dimension 4 rests on.
3. **§ 2.4** — native FHIR repositories, which remove the premise the classic division of labour between the standards rested on.

How to read the table that follows. It is a functional classification, not a scorecard. It compares specified mechanisms and design emphasis rather than assigning unsupported maturity grades.

2.1 Functional Classification of the Standards

LAYER OR FUNCTION	HL7 FHIR	OPENEHR
Principal design emphasis	Standardised representation and API-based exchange, with profiles for particular contexts	Vendor-neutral longitudinal records and clinical modelling through a two-level architecture
Modelling mechanism	Core resources constrained through profiles and extended where necessary	Stable reference model constrained through archetypes and operational templates
Terminology binding	Bindings from coded elements to value sets, with defined binding strengths	Archetype term bindings and external terminology bindings
Standard query mechanisms	FHIR Search and FHIRPath; SQL on FHIR provides a relational projection	AQL over archetype paths, where implemented by the repository
Persistence technology	Not prescribed	Not prescribed
Exchange mechanisms	RESTful API, messaging, documents and bulk-data specifications	REST API and EHR Extract mechanisms, with uneven product support
Clinical modelling assets	Resources, profiles, logical models, implementation guides and Questionnaire/SDC	Reference model, archetypes, templates and curated CKM libraries

LAYER OR FUNCTION	HL7 FHIR	OPENEHR
Secondary-use paths	Bulk Data, SQL on FHIR and transformations into analytical models	AQL extraction and transformations into analytical models

Blobel and Oemig (2018) classify the two differently, by data-model level. Both sit at the **logical** level, they argue, and differ only in view: openEHR archetypes take the informational view, FHIR the computational view.

Two things follow from that reading:

- **The two are not opposites of differing depth.** They are different views at a comparable level, which is why a contest over „modelling depth" misses what separates them.
- **Deep cross-domain interoperability needs a knowledge level**, and neither standard integrates it.

The authors are HL7-affiliated, and their position is not empirically proven. It is carried here as a depolarising grid rather than as a finding.

2.2 The Persistence Statement of the openEHR Source Itself

The tutorial presentation by Pazos, Atalag, Marco-Ruiz, Sundvall and Miranda Freire (2015) carries weight because it comes from inside the openEHR camp. It states plainly that openEHR does not specify persistence, that the choice of database technology depends on context, and that hybrid arrangements are to be recommended.

That supports the central observation of this section: **openEHR is an architectural pattern and a modelling language, not a data format or a database schema.**

The same property cuts both ways:

- **As a strength:** it adapts to the persistence technology the setting calls for, and it does not oblige an organisation to a storage engine.
- **As a weakness:** the interoperability guarantee holds only under disciplined template management, because nothing in the standard supplies it.

This is the point at which the native FHIR CDR gives a different answer.

2.3 The Persistence Statement from the FHIR Camp Itself, and the Technical Case

The FHIR specification is as plain about its purpose as the openEHR tutorial is about persistence, and it says the mirror-image thing. Its overview for architects describes FHIR as "a collection of well-defined interfaces for interoperating between two applications" and states that "FHIR's primary purpose is to address interoperability with well-structured, expressive data models and simple, efficient data exchange mechanisms" (HL7 FHIR R4, *Overview — Architects*; Document 4, Appendix H-2). Exchange is the stated purpose. Persistence is not prescribed, and it is not excluded either: the specification defines resources, interfaces and exchange paradigms and says nothing about how a system holds them (Document 4, Appendix D-1). **FHIR is an interface and content specification, not a database schema and not a modelling language for clinical content.**

The same property cuts both ways here as well:

- **As a strength:** the specification does not oblige an implementer to any store, and a repository may hold resources natively or map them from another model at the interface, which is what § 1.6.1 counts.
- **As a weakness:** the interoperability of two FHIR systems is guaranteed by nothing in the base specification; it depends on the profiles both conform to and on their compatibility (§ 1.5).

The technical case for FHIR rests on four properties that are stated in or evidenced around the specification:

- **Web-native architecture:** REST, JSON, OAuth2/SMART on FHIR. A large tooling landscape: HAPI FHIR, Firely .NET SDK, tested validators, SUSHI/FSH for profile authoring, the Bulk Data Access API.
- **157 resource definitions in R5** cover common clinical and administrative domains, a figure the specification states on its own resource list; resources such as Condition, Observation and Medication provide shared starting points that implementation guides can constrain (Document 4, Appendix H).
- **Integrated publication and conformance framework:** resources, profiles, terminology artefacts and implementation guides use a common specification framework. This does not ensure compatibility between independently governed guides (§ 1.5).
- **Real adoption:** Apple Health Records, Google Healthcare API, the German TI, gematik, MII, NHS England (Genomics, GP Connect), the US Cures Act (ONC § 170.215), the Israeli Ministry of Health and the Australian My Health Record.

That the interface specification has since also become a storage format is the development § 2.4 examines.

2.4 Native FHIR CDRs — the Decisive Technical Lever

> Persistent FHIR repositories can be used for clinical records, which removes the technical premise on which the classic division of labour between the standards rested. Feasibility is documented, comparative effectiveness is not. For a complete primary clinical record held authoritatively in FHIR there is so far only the account of one vendor itself (Document 4, D-7.8); Option C is therefore a demanding solution that has not yet been independently validated.

The most important trend rendering the thesis "FHIR is only an exchange format" obsolete is the maturing of native FHIR Clinical Data Repositories. Here FHIR also becomes the storage level, with the aim of eliminating mapping between persistence and exchange.

What constitutes a native FHIR CDR. Five properties, and the first is the one the others rest on.

- **Data are persisted as they are exchanged** — an `Observation` is stored as a FHIR resource, often as JSONB in PostgreSQL.
- **Create, read, update, delete (CRUD) and history are standardised FHIR operations,** not product-specific interfaces.
- **Versioning (`meta.versionId`, `_history`) and audit trails (`AuditEvent`) are built in,** rather than added beside the store.
- **Terminology services (`$expand`, `$validate-code`, `$translate`) run in the same server,** so validation happens where the data sit.
- **Subscriptions enable event-driven architectures.**

Mature implementations are in use in productive national, regional and research networks, as open source, as commercial products and as hyperscaler PaaS.

What a native FHIR CDR delivers, compared with an openEHR CDR behind a FHIR layer. One plus point is architectural, the other three follow from it.

- **The mapping layer between persistence and exchange disappears.** What remains is the mapping from the source into FHIR, not additionally the mapping from openEHR into FHIR.
- **A standardised route to relational analytics,** through SQL on FHIR where supported.
- **Research handovers** through FHIR Bulk Data (Flat FHIR).

- **Integration with decision-support mechanisms,** including CQL and CDS Hooks, subject to product architecture and implementation.

How far each of the three is productive today differs by tool. § 2.5 records that AQL has been stable for years while SQL on FHIR is not yet its equivalent everywhere.

Production evidence (as of July 2026). The thesis that FHIR-native CDRs are production-ready rests not on product availability alone (Appendix A-1) but on documented production instances, at three different levels and with differing evidence quality:

- **National level:** The Care Data Repository of the Welsh National Data Resource is built as a persistent FHIR store; every clinical and administrative record there exists as a native FHIR resource. Productively implemented are, among others, administrative data, coded drug allergies and paediatric growth charts.
- **Infrastructure level:** The German Medical Informatics Initiative has established FHIR as the implementation basis of the data integration centres of all university hospitals.
- **Institution level:** At University Hospital Essen a relational FHIR store serves as the operational basis for oncological cohort building and for a real-time dashboard in blood-product management that is used daily.

In detail and with evidence grading: Appendix D-7. In summary: the *feasibility* of FHIR-native persistence is demonstrated; its *effectiveness* in the sense of independently measured interoperability gains is not, and the same holds for openEHR.

Limits that must be named in fairness. Profile discipline is a precondition, otherwise "extensionitis" follows, and Appendix D-10 records a member state in which that precondition is documented as unmet by the responsible bodies themselves. Very deep clinical semantics are more laborious to model than in openEHR archetypes. Performance on large historical data volumes requires indexing engineering, which mature products have however solved.

The **storage overhead is substantial**. A study reported in a systematic review (Bennett et al., MIMIC-IV to FHIR) records a considerable increase in storage requirements on conversion to FHIR. The HAPI FHIR format needs markedly more space than the relational source structure (Appendix D-8.2). Workflow orchestration and consent enforcement are not FHIR-native, but the same applies to openEHR.

Added to this is a limitation that Welsh practice itself discloses. Even a FHIR-native CDR does not remove the mapping when the source systems speak HL7 v2: there, the data continue to be transformed from the primary systems into FHIR. What is eliminated is the *additional* transformation stage between persistence and exchange model, not the first between source and target model (cf. § 1.6.2 on the number of transformation layers).

Conclusion. A national or regional infrastructure can use a FHIR CDR to avoid a separate transformation between its persistence and exchange standards, provided that its stored profiles are compatible with the required exchange profiles. Legacy adapters may still convert source data into FHIR. openEHR remains a candidate where its modelling and governance approach provides a demonstrated benefit that justifies the additional FHIR boundary. That benefit is strongest when openEHR governs authoritative clinical capture rather than only a downstream replica (§ 2.7).

The distinction "FHIR exchanges, openEHR stores" has a historical basis and is not an operational decision rule. Both standards reach across it: openEHR specifies mechanisms for extracting records, and FHIR carries persistence-relevant concepts of its own, version history, the `transaction` interaction and audit resources, so one standard family can serve persistence and exchange while profile-to-profile transformations may still be needed. What persists is a difference in modelling philosophy, FHIR's 80/20 heuristic for its core resources with extensions for the rest against openEHR's broad, reusable maximal data set per clinical concept. That difference has practical consequences; it does not establish a structural separation between persistence and exchange. For new

infrastructure the decision is therefore a trade-off between the value of openEHR's clinical-modelling approach on one side and regulatory alignment, ecosystem breadth and fewer cross-standard mappings on the other.

2.4.1 The Three-Domain Thesis and Its Migration into Official Frameworks

In parallel with the conceptual obsolescence of the persistence/exchange distinction (§ 2.4), a related but distinct thesis has established itself. It divides the health data space functionally across three standards: openEHR for clinical documentation and administration, FHIR for exchange, OMOP for longitudinal analysis. This tripartition must be distinguished analytically from the two-layer thesis, because it introduces an additional transformation path.

The thesis is not merely a community position: it stands in official frameworks. The Australian *National Framework for Digital Health Standards* (ADHA, May 2026) assigns FHIR, in its standards taxonomy, to the category *Information Exchange*, and openEHR together with OMOP to the category *Content*. The accompanying figure depicts the two as separate, complementary circles.

The ÖGTelemed position paper (June 2026) adopts in its target picture the separation of persistence and exchange layers and adds OMOP CDM for secondary use. Both documents refer, directly or indirectly, to the same conceptual foundation.

The OECD's 2026 interoperability report provides a further institutional instance. Its illustrative standards example assigns FHIR to exchange and repository integration, openEHR to structured longitudinal records and OMOP to secondary analysis. The report labels the illustration non-prescriptive. Its openEHR description relies on openEHR International and its OMOP role attribution on an industry blog rather than on a comparative architecture study. It therefore establishes that the tripartition has entered international policy discourse, not that the division is technically, clinically or economically preferable (Document 4, Appendix E-10).

Spain provides the newest and most consequential instance. The ministry-commissioned Delphi study of 2026 assigns openEHR to the record and FHIR to the exchange layer for an entire national system, with the mapping between them to be defined in conjunction with the clinical knowledge model (Appendix E-12). It thereby writes the tripartition's missing cost into its own plan: the governed cross-standard boundary of § 1.6 becomes a standing national obligation, accepted in writing.

The substantive critique. A functional division into domains does not solve the central procurement problem. It presupposes what it does not deliver:

- **It assumes the transformations away.** The tripartition takes standardised, low-loss conversion between the layers for granted. Their governance effort and quality assurance are systematically undervalued in the argument.
- **The losses are clinical, not only technical.** Semantic degradation at a layer transition can impair data quality for secondary use and for decision support (§ 1.6.2).
- **It starts where almost no institution stands.** The three-layer model describes a target architecture on the assumption that the primary data already sit in one of the three open standards. For the overwhelming majority of European institutions they do not, and the structural inertia of proprietary HIS systems is the actual obstacle (§ 1.3, Dimension 3).

The provenance of the thesis is worth recording alongside its content. Its most prominent formulation (Tsafnat et al. 2024) comes from institutional representatives of the very standards it recommends, and it is a non-peer-reviewed editorial. A peer-reviewed counter-position is now available (Kapitan et al. 2025). Detailed source criticism appears in Appendix E-1.

A remarkable counter-finding from the openEHR-leaning camp. The ÖGTelemed paper does reproduce the domain separation in its figure, but its body text expressly records that **both openEHR and HL7 FHIR** can be used for the

vendor-independent persistence of clinical data. It names "FHIR as a persistence standard" as a discussion conducted within the HL7 community with real implementations.

A source one would assign to the persistence argument thereby confirms precisely the finding of § 2.4: the dividing line runs not between "storing" and "exchanging" but between vendor-neutral and vendor-bound data holding. That is the analytically more sustainable distinction, and it is achievable with both standards.

2.5 Arguments for openEHR (Technical, Complementary)

> openEHR's two-level architecture, its curated archetype library and its node-level term bindings give clinical modelling a shared, governed home that FHIR distributes across separate mechanisms. **The condition is that none of these is a property of the standard alone:** each depends on the discipline of the body maintaining the models, and the promise that data outlive the application holds only if the operational templates are retained in every version ever used.

- **Two-level modelling.** Clinicians model in a technology-agnostic layer (the archetype), developers implement against a stable reference model: a conceptually clean separation of meaning and technology.
- **A mature clinical model library.** The CKM contains thousands of formally reviewed multilingual archetypes curated over years.
- **Lossless persistence — within the model context.** openEHR is designed to store clinical contexts (observer, method, time series, provenance) without loss of meaning; round-trip fidelity is the standard case within the openEHR ecosystem. This statement must be qualified: it holds **as long as the associated archetype or operational template is available**. Where local at-codes serve as values, the machine-processable meaning is resolvable only against the model (in detail § 2.5.1). Losslessness is thus a property of the system of data *and* models, not of the data alone.
- **AQL is powerful** and has been stable for years; SQL on FHIR provides a comparable approach for projecting FHIR into queryable tables, but it is not yet consistently production-ready or broadly supported across all FHIR platforms.

2.5.1 Self-Description of Instance Data — openEHR's Model Dependency

A point rarely argued with precision in the debate concerns the question of whether persisted openEHR data are **interpretable without the associated archetype**. It matters considerably for the procurement decision, because openEHR's central promise, "the data outlive the application", depends on it.

What the specification says. The openEHR reference model knows, for `DV_CODED_TEXT`, the constant value `local` as terminology identifier. The Data Types specification records that this value means the origin of the permissible values is described **within the archetype**.

What is meant are the so-called **at-codes** (from *archetype term*) — archetype-internal placeholder codes of the form `at0001`, `at0006`. They perform two functions. First, they **identify nodes**: an archetype path is nothing other than the alternating pattern of reference-model attribute names and node codes. Second — and this is the critical case — they serve **as values**.

The Architecture Overview illustrates this with the ORDINAL example. An ELEMENT node identified by `at0005` can take the values 0, 1 and 2, which are in turn coded by `at0006`, `at0007` and `at0008`. The associated rubrics ("mild impairment" and so forth) are located in the ontology section of the archetype, not in the instance.

Binding internal archetype terminology to external terminologies is an additional, optional layer. The openEHR specification cautions that external terminologies often have incomplete coverage and that such mappings may be approximate, so they require careful curation.

What that means concretely. In canonical openEHR representation, a `DV_CODED_TEXT` instance carries both the rubric in `value` and its `defining_code`, a terminology identifier plus code. The instance-level coding is therefore present in the data, and coded values remain human-readable. It would be inaccurate to claim categorically that coding decisions are absent from the data; rather, the archetype or template supplies essential constraints and contextual semantics in addition to the instance data.

In Better/EHRbase-style FLAT JSON, a local archetype code may be represented by the code alone, because its rubric and terminology identifier are recoverable from the applicable template. This shortcut does not generally apply to externally coded values, for which code, terminology and value must be supplied. The operational template (`.opt`) is the compiled artefact used for operational, template-based validation.

The robust finding is therefore narrower and at the same time sharper. Where local at-codes are used as values, the *machine-processable* meaning is resolvable only against the archetype, and it is that meaning which makes the value usable for secondary use, decision support and analytics. The rubric text is text; the semantics live in the model.

Two senses of "self-describing", and only one of them is contested. *Instance-readable* means that a person can read a coded value out of the instance: an openEHR `DV_CODED_TEXT` carries a rubric alongside the code, so no terminology service is needed to see what was meant. *Semantically self-sufficient* means that the instance alone carries everything required to interpret and validate it: the constraints, the path semantics, the value-set definitions and stable global term identities.

For canonical openEHR instances the first is often true. The second is generally not true where local at-codes are used, because their semantics are defined in the archetype or the operational template. The same discipline applies to FHIR: `system` plus `code` gives a portable terminology identity, while the resource and its profile determine how that coding is meant to be used. What follows concerns the second sense throughout.

The comparison with FHIR. In the coded case, a FHIR `CodeableConcept` contains one or more `Coding` elements, each of which can carry a `system`, an absolute URI identifying the code system, alongside a `code`. The base specification does not universally require this: `Coding.system` is 0..1, and a `CodeableConcept` may be represented by free text alone. Where `system` and `code` are supplied, however, the coded identifier is portable beyond the resource's local structural context, even though the system URI need not be dereferenceable.

By contrast, an openEHR at-code is local archetype terminology: its meaning is determined by the terminology section of the applicable archetype and, in operational use, by the template context. An at-code on its own is therefore not an independently meaningful global identifier.

That qualifies the claim that openEHR persistence is self-describing, and does not disprove it. Canonical instances may carry human-readable rubrics and codes, but complete interpretation of locally coded content depends on the associated archetype or template. FHIR for its part is more than a transport format, because its resources, datatype definitions, profiles, terminology bindings and extensions form a persistent semantic model.

That is a real structural difference, and it runs counter to the widespread narrative that openEHR persistence is per se self-describing while FHIR is merely a "transport format".

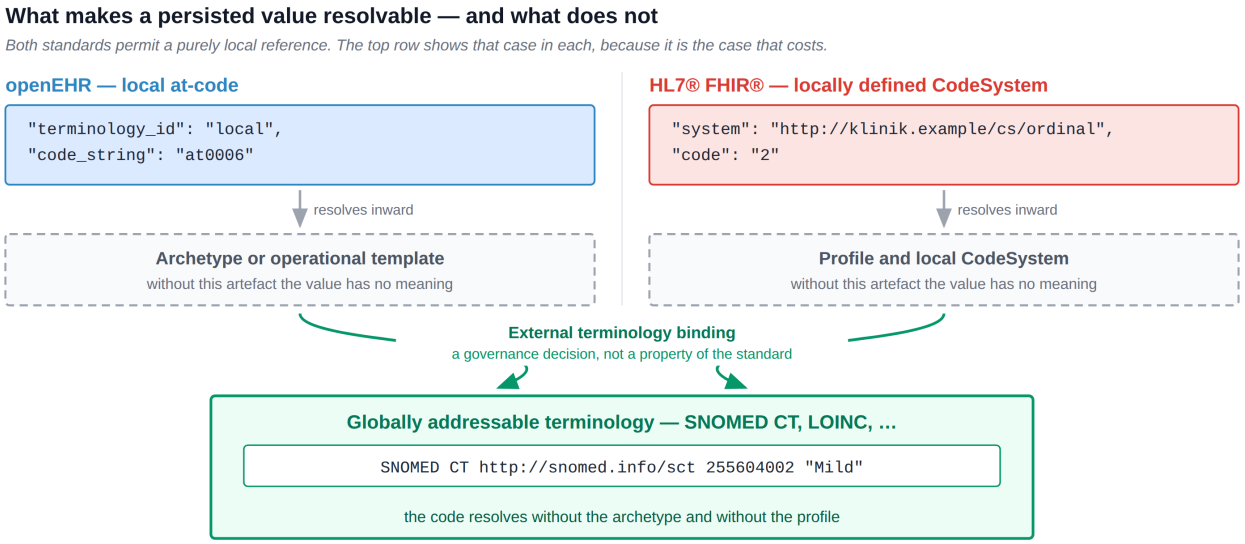

Figure 6 — Both standards permit a purely local reference, and the figure shows that case in each, because it is the case that costs: the value then resolves only against the artefact that defines it. What lifts it on either side is external terminology binding — a governance decision of the modelling organisation rather than a property of the standard.

An empirical addendum on practice. The external terminology binding whose absence is described above as a risk is common practice in the FHIR research environment. In the systematic review by Vorisek et al. (2022), 29% of the included studies report using SNOMED CT and 37% LOINC (Appendix D-8.3). A methodologically comparable inventory of the terminology-binding rate of published openEHR implementations does not exist. This is not evidence of better semantics, but of **documented binding practice** — precisely the property on which the resolvability of persisted values depends.

An honest delimitation: the point is gradual, not categorical.

FHIR is not free of model dependency either. A profile determines which extensions are to be interpreted how, and locally defined `CodeSystem` resources raise the same resolution question.

On the openEHR side the gap is remediable in part. An archetype consistently bound to external terminology such as SNOMED CT or LOINC makes the clinical *meaning* of a coded value resolvable without the model, precisely the deficit the at-code creates.

In the terms set out above, that buys instance-readability. It does not buy semantic self-sufficiency. Structure, cardinalities, permissible-value sets and the version context still resolve only against the archetype or the operational template. Terminology binding therefore does not lift the retention obligation set out in the next paragraph. It is the same qualification § 2.5 states for losslessness, which is a property of the system of data *and* models rather than of the data alone.

Two questions remain outside the standard itself. Whether the binding happens at all is a **governance question for the modelling organisation**. Whether anything obliges a consumer to honour it is a question of binding strength, treated at § 2.5.2.

Consequence for the architecture and procurement decision. The promise "the data outlive the application" holds only if the **model repository also survives**. Whoever procures openEHR persistence implicitly procures an obligation. The operational templates have to be retained in every version ever used productively, for as long as the data themselves are retained, which for clinical records is measured in decades.

This requirement is formulated nowhere in the recommendations evaluated, including the Gartner Research Note (Appendix D-6) and the ÖGTelemed position paper (Appendix E-3). It must be carried as a lock-in dimension in its own right (§ 1.3, Dimension 4).

2.5.2 Binding Strength — Where the Difference Actually Sits

§ 2.5.1 leaves one question open, and it is the one that decides the case. On both sides the outward pointer is optional: `Coding.system` is 0..1 in FHIR, and the `term_bindings` section is optional in openEHR. What decides whether an instance can be resolved is therefore not which standard was chosen but whether the binding was specified, and whether anything checks it. In the terms set out in § 2.5.1, this is where the terminology part of the second sense is decided. Binding strength changes nothing about what a person can read out of an instance; it changes what a validator can require of it.

FHIR makes the degree of obligation explicit and machine-readable. `ElementDefinition.binding.strength` carries one of four values, and only one of them compels anything.

Required: "To be conformant, the concept in this element SHALL be from the specified value set." *Extensible*: the same obligation, but only where a code in the value set can apply to the concept at all. *Preferred*: instances are encouraged to draw from the codes, and remain conformant if they do not. *Example*: instances are "not expected or even encouraged" to draw from the value set, which merely illustrates the kind of concept intended. **Evidence grade A** — this is what the normative specification requires of a conformant system.

The consequence is uncomfortable and belongs beside the at-code finding. **A profile whose safety-critical elements carry `example` bindings is fully conformant FHIR and enforces nothing.** The pointer leads outwards only where a profile makes it lead outwards, and the four strengths are precisely the instrument that records whether it does. Where the strength is not stated, the outward reach of § 2.5.1 is a possibility rather than a property.

Even a `required` binding buys only part of the second sense. It fixes the terminology of an element, not the structure, the paths, the cardinalities or the version context; those continue to resolve against the profile.

Read from the tooling rather than from the specification, the gradation collapses further. The four strengths describe what a conformant system is obliged to do. What a validator can actually decide is narrower.

An interoperability analysis of two implementation guides records the practical position (Kramer & Moesel 2023; Appendix D-14):

- **`preferred` and `example`** do not constrain the permitted values at all.
- **`extensible`** permits an additional code wherever the concept is not covered by the value set. No algorithm can make that judgement without semantic understanding, so tools fall back on accepting any code.
- **`required`** is consequently the only strength that definitively limits the set of instances a profile admits.

The same paper shows what that costs in a concrete exchange. US Core binds `communication.language` extensibly, IPS binds it as required. The US Core value set is a *subset* of the IPS one, and the exchange is still classified potentially infeasible. The reason is that the extensible binding on the sending side lets a code through that the receiver must reject.

The procurement consequence is one step stronger than the specification alone supports. For safety-critical content, `extensible` should be read as an unbound element with a recommendation attached, not as a weaker form of `required`.

openEHR has the mechanism but, so far, not the gradation. ADL and the Archetype Object Model define a `term_bindings` sub-section (ADL2 specification, § 7.13.4; AOM2 `ARCHETYPE_TERMINOLOGY.term_bindings`) in which at-codes and id-codes are mapped to codes from external terminologies such as SNOMED CT. Populating it is

optional, and the specification states no obligation to do so. What the section does not carry is a binding strength: a binding is present or absent, and the archetype cannot express whether a consumer must honour it.

That gap is recognised inside openEHR, and the proposed remedy is the one in FHIR. On 19 August 2020 Thomas Beale, for the openEHR Specifications Editorial Committee, proposed adding `required | extensible | preferred | example` to AOM and ADL. He named "the model used in HL7 FHIR" as the template. The motivation he cites is the rigidity of the previously implicit `required` semantics. The thread runs into 2023 without a visible conclusion.

Evidence grade D — a community forum — weighted higher as a statement against the interest of the camp that makes it, on the same reasoning applied at § 1.3, Dimension 1. It is carried here as a convergence signal and as corroboration that the gradation is regarded as necessary, not as evidence that it exists.

What follows for procurement, and what does not. Three things follow, and they are conditions rather than a recommendation.

1. **For safety-critical elements, ask for the strength, not only for the binding.** For medication and allergy content, a binding at `required` against a published, versioned value set is the only strength that a validator can enforce. `example` on such an element should be treated as a defect regardless of the standard chosen, and the same demand applies to the term bindings of an openEHR operational template, where it must be stated contractually because the model cannot yet express it.
2. **Ask where the value set is resolved and by what.** A binding is worth what the resolution point is worth: a published, versioned value set and a terminology server that serves it, rather than a code list inside an implementation.
3. **Ask what checks it, and when.** A binding that no validator exercises in the release pipeline is a declaration of intent. This is where the requirement returns to the conformance-testing argument of § 1.5 and Part I, Finding 4.

Semantic annotation is self-discipline, and both standards permit it, but not by the same mechanism and not with the same reach.

Naming a concept is not annotating it. A model node labelled "body height" is readable and still local. The same node annotated with SNOMED CT `1153637007 |Body height (observable entity)|` carries an identity that survives the model it was written in. Both standards make that possible and neither compels it.

- **openEHR attaches the annotation to the node itself.** The optional `term_bindings` section maps an at-code or id-code, the archetype's own node identifier, to an external concept. Any node in an archetype can therefore be given an external identity (§ 2.5.1).
- **FHIR keeps the two things apart.** A `binding` attaches to an element of a coded type and constrains the *values* that element may carry. It is not an annotation of the node, and it cannot sit on an element that carries no code.
- **Annotating the definition is a third mechanism in FHIR.** `ElementDefinition.code` attaches concept codes to the element definition, and is used in logical models as well as profiles.
- **Forms add a fourth choice.** `answerOption` enumerates permitted answers inside the form and leaves them semantically local; `answerValueSet` binds the item to a published, resolvable value set. `item.code` gives the question itself a coded identity.

The inequality at the node layer is real, and it runs in openEHR's favour. One optional section reaches every node of an archetype, where FHIR distributes the same capability across binding, `ElementDefinition.code` and the Questionnaire mechanisms, each with its own scope. What follows is a difference in how easily the discipline can be exercised and audited, not a difference in whether it is possible.

For the comparison, the conclusion is that annotation discipline is a property of neither standard. It is a governance property of the body that writes the models. How unevenly it is exercised even inside one ecosystem is

counted below (§ 2.5.2). The desideratum that follows belongs to whoever maintains a model holding, on either side, and it is stated in Document 6 rather than resolved here.

What the German corpus actually does, and why it cuts both ways. The propositions above concern what the two specifications permit. Deployed practice was counted separately, at Appendix D-10.7, and the count corrects the expectation this section raises.

Across 2,996 StructureDefinitions the German families declare 2,399 bindings of their own. Of those, 79.4% are `required` and 1.9% `example`. **Where this corpus binds, it binds hard, and the `example` risk described above does not materialise broadly.** Two qualifications of that count stand in Document 5 (D-10.7): a binding census understates the rigour that fixed and pattern constraints add, and the safety-critical medication and allergy profiles are the weakest part of the corpus rather than the strongest. The first of the three conditions above therefore stands undiminished for medication and allergy content: ask for the strength, not only for the binding.

2.6 Technical Points of Contention with Truth on Both Sides

The standing objections raised against openEHR in the HL7 materials are neither adopted nor dismissed as a group. Measured one at a time against the specifications, most turn out to be half right and to compare layers that do not correspond. **One of them stands apart, and it is the one to read first: § 2.6.3, the objection that carries independent empirical support** — the difficulty of the openEHR modelling process at scale, corroborated by Christensen & Ellingsen for the FIKS project. The template-versus-questionnaire debate (§ 2.6.1) and the storage, architecture and tooling objections (§ 2.6.2) are argument against argument. And the counts available for the German FHIR corpus have no openEHR equivalent to be compared against.

Where these objections come from, and what is done with them here. The HL7-side materials assessed for this document carry a set of standing objections to openEHR. They circulate in the debate as though they were settled. Each is stated below as it is actually made, then measured against the specifications and the sources. In most cases what remains is a claim that is partly right, and that is why they persist: an objection with a true core outlives one without.

Naming that core, and naming the point at which it stops, is more useful for a procurement decision than deciding who wins the exchange. This section has two counterparts. § 1.4 gives the arguments of the openEHR side the same hearing, and Appendix E of Document 4 examines the recent position papers of both camps for what they can carry.

2.6.1 openEHR Template versus FHIR Questionnaire/SDC

> *“ openEHR covers FHIR's Questionnaire/QuestionnaireResponse equivalents only*

Accurate about the form, wrong about the model. The derivation behind this verdict, with the query mechanics, the survey figures and the review figures, is at Appendix D-15 of Document 4; this section carries its conclusions.

The objection compares two different layers, and that is why it half holds. openEHR carries a reference model, archetypes and templates; FHIR carries resources, profiles or logical models, and Questionnaires. The form layers correspond, template and Questionnaire, and so do the model layers, archetype and profile or logical model. Read at the form layer the objection is essentially right: Questionnaires with Structured Data Capture (SDC) deliver nested groups, terminology binding, conditional logic, modular reuse and calculated structures, so the argument "FHIR alone cannot represent this depth" is not tenable. Read as a claim about the whole of openEHR it is a category error, and the FHIR community's own practice says so: where an archetype was published beside a FHIR artefact, the counterpart carried beside it was the Logical Model, not a Questionnaire (Appendix E-7).

Three differences survive the comparison, and none of them is about expressive power.

First, the query key. AQL selects a captured value by its archetype path, which every template embedding that archetype shares. A `QuestionnaireResponse` item is addressed by a `linkId` that is local to one form, and the releases the European and German profiles are pinned to define no search parameter for it. FHIR reaches the same place by one further step, SDC extraction into an ordinary resource, and that step is itself a mapping with the obligations of § 1.6; composite search parameters in the current build will close the gap at the source once servers implement them. The advantage is bounded by products rather than by the language: in the vendor survey the majority of openEHR repositories declare AQL, not all of them (Appendix A-2).

Second, where the shared model layer lives. On the FHIR profile layer reuse is machine-readable through `baseDefinition`, and for the German corpus the count exists in Document 5 (D-10.7): reuse is real and local, and the shared national layer is not the structural parent of the landscape. On the openEHR side the equivalent is measurable in principle but unpublished. On the Questionnaire layer SDC specifies composition, `subQuestionnaire` and `$assemble`, but establishes no curated library of form modules. Such a library is part of the standard in openEHR and something a community builds beside it in FHIR, and the REDCap Shared Library shows that it can be built.

Third, the reach of the instance's pointer. An at-code keeps its identity across every form that embeds its archetype; a `linkId` is scoped to one Questionnaire. Neither is self-sufficient without its definition artefact, but one is portable between forms and the other is not.

A fourth observation belongs to forms rather than to either standard. Form-driven capture often asks for derived, purpose-specific facts that no ward documents in that form, and deriving them is itself a mapping with the obligations of § 1.6. The German emergency-department network AKTIN builds its register reporting on exactly that principle, with no shadow documentation beside care and Questionnaires pre-populated from the record (Appendix D-13). A Questionnaire is therefore not automatically a form for capture at the point of care.

Modelling depth does not separate the two standards. Three systematic reviews of differing methodology document FHIR in oncology, genomics, infectious disease and imaging, used predominantly for standardisation and primary capture rather than for exchange (Appendix D-8). That removes one argument from the debate: the claim that FHIR is structurally too shallow for research-grade modelling does not hold as a blanket statement. It does **not** show FHIR superior in modelling depth, because none of the reviews includes openEHR as a comparator and none measures effectiveness.

The difference is therefore not "can or cannot" but design intent, maturity and where the governance effort falls. openEHR supplies a shared clinical model library and a query key that is shared with it. FHIR supplies a larger and more distributed ecosystem in which the same properties are assembled per guide. For the architecture decision the choice between openEHR template and Questionnaire with SDC is a trade-off between supplied clinical governance and ecosystem uniformity, not a question of fundamental feasibility.

2.6.2 Storage, Architecture and Tooling — Three Further Objections

> *“ openEHR's internal persistence format is proprietary*

Partly accurate, but gradual. The openEHR reference model is open. Concrete CDR storage layouts are vendor-specific, and the same holds for FHIR servers. The more robust objection lies elsewhere and is usually missed in the debate: not in the storage layout but in the **model dependency of the instance data**. Where local at-codes serve as values, the machine-processable meaning is resolvable only against the archetype (§ 2.5.1). That is not a proprietariness problem but a self-description problem, and it strikes the openEHR argumentation at a more sensitive spot than the usual accusation.

More recent measurements move the objection from the storage layout to the exit route. A peer-reviewed survey put a fifty-four-question instrument to the vendors of seventeen production-capable openEHR repositories; eleven answered (Delussu et al. 2024; Appendix A-2). What the answers show is not a proprietary storage layout, which was never the interesting part, but the state of the two mechanisms the specification provides for getting content out.

The **openEHR EHR Extract**, which exists for precisely that purpose, is declared by **one of the eleven**, with a second offering it through a separate product. Composition serialisation is dispersed across canonical XML, canonical JSON, FLAT, STRUCTURED, a vendor Template Data Document and a vendor AQL-path/value JSON, and reference-model versions run from 1.0.1 to 1.1.0 with no common floor.

The authors build openEHR tooling themselves, and they draw the conclusion in their own words: those differences prevent straightforward data interoperability between the repositories unless an external intervention is applied. What is missing to close the gap is a conformance testing framework, the openEHR conformance guide being still in development.

The objection as usually stated is therefore too narrow rather than wrong. The storage layout is nobody's problem. The absence of an implemented, tested exit route is. It is the same class of problem as the one Document 5 counts on the FHIR side: not a defect of the specification, but of what the products built from it were held to.

> *“ No reference architecture, only a generic information model*

Factually accurate for the classic openEHR information model. The most recent ISO/EN 13606 revision does contain, for the first time, a reference architecture including a FHIR mapping option. In productive openEHR practice this is not yet consistently visible.

> *“ Queries can only be made with a special editor*

Correct in form, not in substance, and the substance is narrower than the objection. AQL is mature; the tooling is smaller and more proprietary than the FHIR ecosystem. What the objection misses is that the editor is not the constraint. In the vendor survey above, **AQL is declared by six of the eleven repositories**; four are queried in SQL, one of them alongside AQL, and two through a query formalism of the vendor's own (Appendix A-2).

AQL adoption is rising: the 2013 ratio of five AQL to eight SQL has reversed, and the direction matters. But an argument that treats AQL as the portable query key of the openEHR ecosystem holds for the products that implement it, which in 2023 was a bare majority. That is a finding about products rather than about the language, and it is recorded that way.

2.6.3 One of These Objections Has More Than Argument Behind It

The four entries above are contested claims weighed against specifications and sources. One of them additionally has independent empirical support. It is recorded separately for that reason, rather than left in the list where it would read as one objection among others.

The objections to openEHR modelling practice voiced in the HL7 materials are corroborated by Christensen & Ellingsen (2016), the only peer-reviewed long-term case study of a large-scale openEHR implementation (Appendix D-5). That corroborates the difficulty of the modelling process at scale. It does not corroborate the other three objections, and the study itself argues for better governance rather than against the standard.[2]

[2] **On the apparent double role of this study.** Christensen & Ellingsen (2016) is invoked at several points with differing thrusts. It serves as the methodologically most robust evidence *against* an unconsidered large-scale openEHR introduction (§ 1.6, this § 2.6, Appendix C-4 point 2). It also carries the note that the failure of the assumptions examined is *not* evidence of a failure of openEHR as an idea (Appendix D-5). Both readings

2.7 OMOP as a Third Standard — the Research Perspective

Oemig et al. (2018) treat, alongside FHIR and openEHR, the Observational Medical Outcomes Partnership (OMOP) Common Data Model of the OHDSI initiative. OMOP is a fixed relational schema combined with standardised vocabularies, mapping infrastructure and a library of reusable analytical methods. That combination, not the schema alone, is what has made it the most widely adopted harmonisation model for observational data.

The claim that OMOP is the clinical data model, and what the sources support. A view circulates in the debate that OMOP is the clinical data model proper and stands in competition with openEHR for that role. It deserves a hearing, because a part of it is supported.

Finster et al. (2025) review the common data models and data standards available for tabular health data. They assess i2b2, the Sentinel CDM, the PCORnet CDM and OMOP against CDA, HL7 v2, FHIR and openEHR, across suitability, popularity, adaptability, interoperability and support, and find OMOP and FHIR the highest scoring overall. Whoever says that OMOP is established, well supported and more than a warehouse schema is therefore right, and can cite a peer-reviewed review for it.

How much of the literature actually uses it, counted. Across the 99 studies of the 2026 systematic review, OMOP-CDM is the most frequently used of the three standards: 56 studies against 39 for FHIR and 8 for openEHR. Of all 99, 91% applied their standard to a secondary purpose rather than to direct care (Marfoglia et al. 2026; Appendix E-9).

Read carefully, that count supports the reading above rather than the coronation. OMOP-CDM leads the *research* literature, which is the setting 87% of the studies come from, and the review assigns it the analytical role while assigning persistence elsewhere. It measures publications, not installed systems, and the two should not be confused.

Where the claim stops. The same review concludes that no single global representation can be selected and that the ability to transform between representations is what matters, which is the opposite of a coronation. Its subject is tabular health data and its criteria include popularity and support, so a high score is in part a measure of adoption rather than of modelling capability.

OHDSI's own definition places the model where its design intent lies. It is an open community data standard "designed to standardise the structure and content of observational data and to enable efficient analyses that can produce reliable evidence". Its named purposes are characterisation, population-level effect estimation and patient-level prediction.

No source examined for this assessment documents OMOP in use as the system of record for primary clinical documentation, in the sense of § 1.6.1. Absence of documentation is not proof of absence, and it is recorded here as the state of the search rather than as a finding about what is possible.

The two are therefore mostly not competitors but neighbours in different layers. openEHR can capture and persist clinically structured primary data over the long term. OMOP can derive from it a projection harmonised for research and real-world-evidence analysis, as it can from FHIR, CDA, proprietary or claims sources. The transformation is professionally valuable and it is not neutral: mapping rules and any semantic losses have to be documented transparently and validated, under the same requirements as any other mapping (§ 1.6.2).

---

are correct. The authors refute four specific assumptions of model-driven development: a clean separation of clinical and technical concerns; clinical self-governance; broad user commitment; and modelling as a simple mapping of existing practice. Those are assumptions about the *implementation logic at large scale*, not about the sustainability of the standard. The authors plead expressly for formalised modelling structures inside the hospital organisation, not for abandoning openEHR. The study therefore carries two statements from one finding: introduction costs and risks are real and were substantially underestimated in the case documented, and they are addressable given appropriate governance.

The route into OMOP is not an openEHR speciality, and the tooling runs in both directions. Transformation between FHIR and OMOP is an established field of work, with published implementations on both legs:

- **Lenert et al. (2021)** describe a repository that holds clinical data as FHIR as its canonical model and produces PCORnet and OMOP research marts from it automatically, using FHIR Subscription to keep the delay to a minimum. That is the arrangement this document recommends for new infrastructure, with the research projection sitting downstream of the store rather than beside it.
- **Jayathissa et al. (2025)** describe the opposite leg, an ETL from FHIR Bundles into the OMOP CDM.
- **Essaid et al. (2024)** run the direction that is easily forgotten: transforming an existing OMOP warehouse into US-Core-conformant FHIR resources and exporting them through Bulk FHIR, with a non-conformance rate below one percent.

OMOP is therefore a source as well as a target, and the crossing between it and FHIR is as much a governed model boundary as the one between openEHR and FHIR.

That the projection loses meaning is not an inference here but a published finding. Kohler et al. (2023) built the openEHR-to-OMOP path deliberately: the OMOP Conversion Language OMOCL for declarative archetype-to-CDM mappings, and the Eos tool that executes them on real-world and sample records. They record as their own conclusion that transforming openEHR data into the less expressive OMOP CDM leads to a loss of semantics. They also name a weakness on the source side, openEHR's unspecific demographics model, for which they propose two remedies.

A statement of loss from the group building the bridge carries more than the same statement from either camp. That is why the loss is asserted here as a property of the path rather than as a suspicion about the target model.

The symmetry has to be stated, because the evidence is uneven rather than one-sided. Semantic loss on projection into OMOP is documented for the openEHR path because someone examined that path for it. The FHIR-to-OMOP work cited above reports feasibility, conformance rates and production delay. One measurement does exist on that side, but it is thin. The journal version of Rinaldi & Thun 2021 scores both mappings against ISO/TR 12300: equivalence 0 for openEHR to FHIR, a complete match that needed no extension, against 0.9 for FHIR to OMOP, with a good match for about 78% of the elements and the gaps named — DiagnosticReport and ServiceRequest have no counterpart in OMOP, and commentary fields go into a separate Note table. It is a self-assessment by the mapping team, which the authors call subjective and which they score themselves at 2 for method of validation and at 2 for documentation of decisions, on one microbiology dataset (Appendix D-8.3, evidence grade C). It narrows the asymmetry and does not remove it.

Absence of a loss measurement is not a finding of no loss. The target schema is fixed and vocabulary-bound whatever the source, so a source-model detail with no place in it is dropped on every route. What differs between the architectures is therefore not whether the OMOP projection costs something. It is how many governed boundaries the content has already crossed before it gets there, which is the count of § 1.6.1 and not a property of OMOP.

What follows for a target architecture. Whoever designs a national architecture should understand FHIR, openEHR and OMOP as different and in part complementary building blocks, while observing the delimitation of § 2.4.1. The three-domain thesis criticised there asserts a fixed functional assignment (openEHR stores, FHIR exchanges, OMOP analyses) and suppresses the transformation costs between the layers.

The position taken here is narrower. OMOP is a legitimate target format for defined research use cases, into which data are projected from the primary holdings, for instance via FHIR Bulk Data export or AQL extraction. It is not a third mandatory layer of a target architecture, and every projection into the fixed schema is a lossy transformation subject to the same validation and governance obligations as any other mapping.

That the projection is increasingly practised in European research networks such as EHDEN and the German Medical Informatics Initiative demonstrates its usefulness for the research purpose. It does not demonstrate the correctness of a general domain tripartition.

What would change this classification. One finding would move the model out of the research layer and into the comparison this document conducts between openEHR and FHIR. That finding would be a documented instance of OMOP holding the clinical record authoritatively for care delivery, rather than receiving a projection of it. Nothing found for this assessment does that.

Appendix F: Glossary of Abbreviations and Terms

Included are abbreviations and technical terms used in this set (Documents 3 to 5) without introduction. Not included are everyday abbreviations (EU, US, UK), company names and bibliographic notations (DOI, ISBN). Where a term is treated in detail in the text, the section is indicated.

F-1 The Central Terms

These five suffice for reading Part I:

TERM	MEANING
CDR	*Clinical Data Repository* — a repository that retains clinical data for operational or secondary purposes. Its authoritative model and relationship to source systems are central to this assessment.
HIS	Hospital information system (German: *KIS*) — a hospital's central clinical and administrative software. Many established products retain data in internal models and expose standards through interfaces (§ 1.3, Dimension 3).
TCO	*Total Cost of Ownership* — acquisition, operation, maintenance, staffing, migration and exit costs over a system's lifetime (§ 1.6, Appendix C-5).
Mapping	A governed relationship or transformation between data models. It creates validation and maintenance obligations and may lose meaning; its number is an architectural indicator rather than a measure of cost or safety (§ 1.6, § 1.6.2).
EHDS	*European Health Data Space* — EU Regulation (EU) 2025/327. Its technical implementation work is a major driver of European FHIR exchange specifications; the Regulation itself does not prescribe an internal persistence model.

F-2 Standards, Specifications and Data Models

ABBR.	EXPANSION / EXPLANATION
HL7 FHIR	*Fast Healthcare Interoperability Resources* — an HL7 International standard defining resource-based representations, APIs and exchange mechanisms. R4 and R5 are prominent published releases in the material assessed here.
openEHR	Open standard for structured, long-term clinical data holding. Two-level model of a stable reference model and clinically defined archetypes.
Archetype	A reusable model of a single clinical concept (blood pressure, allergy) in openEHR. Designed as a *maximal model*: represents the subject matter as completely as possible.
Template	A combination of several archetypes for a concrete use case (for example an admission form). Constrains the archetypes to what is needed.

ABBR.	EXPANSION / EXPLANATION
Operational template (.opt)	The technically deployed, fully resolved version of a template. Authoritative for validation — and for the interpretability of stored data (§ 1.3 Dim. 4, § 2.5.1).
at-code	Archetype-internal placeholder code (`at0001`, `at0006`, from *archetype term*). Designates structural nodes and in part serves **as a value**. Its machine-processable meaning resides in the archetype, not in the data (§ 2.5.1).
RM	*Reference Model* — openEHR's technical reference model: defines which structures exist (COMPOSITION, SECTION, ENTRY, CLUSTER, ELEMENT), not their clinical meaning.
AQL	*Archetype Query Language* — openEHR's SQL-like query language, which queries along archetype paths.
ADL / TDL	*Archetype Definition Language* / *Template Definition Language* — formal languages used to define archetypes and templates. Their use requires specialised modelling skills (Appendix C-4).
CDA	*Clinical Document Architecture* — an older, document-centric HL7 standard (discharge letters, reports). The basis of many national legacy infrastructures (including ELGA).
HL7 v2	Message-based HL7 legacy standard, still the de facto wiring of many hospitals.
OMOP (CDM)	*Observational Medical Outcomes Partnership Common Data Model* — a fixed relational schema for research analyses. A target format for secondary use, not a persistence or exchange layer (§ 2.7).
Profile	In FHIR: a constraint/refinement of a resource for a particular context (for example AU Core, the German base profiles). Profiles are where FHIR interoperability actually arises — or fails (§ 1.5).
Extension	In FHIR: an extension by fields the base resource does not provide. Uncontrolled proliferation ("extensionitis") undermines interoperability.
Logical Model (LM)	In FHIR: a domain model expressed as a `StructureDefinition` with `kind = logical`, describing which data a clinical field comprises, independently of the resources that later carry them. The subject of Appendix E-7.
DAM	*Domain Analysis Model* — a domain-level analysis model of a clinical field (Appendix E-7).
SDC	*Structured Data Capture* — a FHIR implementation guide for structured form-based capture, built on `Questionnaire`. Central to the modelling-depth debate (§ 2.6.1).
Questionnaire / QuestionnaireResponse	FHIR resources for form definition and form response.
StructureMap	FHIR mechanism for data transformation, inter alia for converting form responses into clinical resources.
CodeableConcept	FHIR data type for coded values. Carries a `code` with, in the normal case, a `system` URI, so that the pointer to the meaning is globally resolvable. `Coding.system` is formally optional (§ 2.5.1).

ABBR.	EXPANSION / EXPLANATION
DV_CODED_TEXT	openEHR data type for coded values. Can point to external terminologies **or** to archetype-internal at-codes (§ 2.5.1).
CQL	*Clinical Quality Language* — a language for clinical rules and evaluation logic in the FHIR environment.
FHIRPath	A path language for navigating FHIR resources.
SQL on FHIR	A specification for relational querying of FHIR data — the FHIR counterpart to AQL.
Bulk Data	FHIR specification for mass export of data (research, analytics).
SMART on FHIR	A standard for connecting apps to FHIR systems with OAuth2-based authorisation.
IPS	*International Patient Summary* — the HL7 FHIR implementation guide for a cross-border summary record; the receiving profile set in the Kramer & Moesel analysis (§ 1.5, § 2.5.2; Appendix D-14).
US Core	The United States national FHIR base implementation guide, referenced by the US certification programme; the sending profile set in the same analysis (§ 1.5; Appendix D-14).
IHE	*Integrating the Healthcare Enterprise* — an organisation defining integration profiles (XDS, MHD, PIXm, PDQm) for concrete use cases.
ISO 13606 / EN 13606	ISO/CEN standard for EHR communication, historically derived from openEHR. A possible governance connection point for Europe (§ 1.4, § 2.6.2).
System of record	The system whose version of a clinical fact is the authoritative one, against which every other copy is reconciled. Distinct from any system that merely stores or displays the same content.
Semantic annotation	Attaching an externally defined, resolvable concept identity to a model node or a question — a SNOMED CT, LOINC or ATC code — rather than describing it in a label alone. Possible in openEHR through term bindings and in FHIR through terminology bindings, `item.code` and `answerValueSet`; compelled by neither (§ 2.5.2).

F-3 Terminologies and Classifications

ABBR.	EXPANSION
SNOMED CT	*Systematized Nomenclature of Medicine — Clinical Terms.* A comprehensive clinical terminology; the most important reference point for semantic unambiguity.
LOINC	*Logical Observation Identifiers Names and Codes.* Coding of laboratory values and observations.
ICD	*International Classification of Diseases.* The WHO classification, primarily for billing and statistics — not a clinical data model.
DRG	*Diagnosis Related Groups.* A case-based flat-rate system built on ICD coding.
ValueSet / CodeSystem	FHIR resources for defining permissible value sets and custom coding systems.

F-4 Architecture, Operations and Procurement

ABBR.	EXPANSION / EXPLANATION
ACID	*Atomicity, Consistency, Isolation, Durability* — the four properties a database write must have if a clinical record is to be relied on: a multi-part write happens completely or not at all; no write leaves the store in a state its own rules forbid; concurrent writes do not read each other half-finished; what has been acknowledged survives a crash. A repository that holds the authoritative record needs all four; an exchange endpoint in front of a proprietary store does not supply them.
TCO	*Total Cost of Ownership* — see F-1.
ETL	*Extract, Transform, Load* — a data transfer path from source systems into a target system. Every ETL path is a transformation layer.
Transaction	The word carries four distinct meanings in this field and is therefore avoided as a general term in this set. **IHE**: a named, specified message exchange between two actors in an integration profile, cited by code (ITI-41, ITI-18); used in that sense in Document 5. **FHIR**: the `transaction` interaction, a bundle applied atomically; a persistence-relevant capability, § 2.4. **Database**: transactional behaviour in the ACID sense, one of the conditions on an authoritative record, § 1.6.3. **ART-DECOR**: a use-case-specific derivation of a dataset that carries cardinality where the dataset does not, Appendix E-7. Where this set means the exchange use cases it examines, it says **data exchanges**, not transactions.
API	*Application Programming Interface.*
REST	An architectural style for web interfaces; the technical foundation of FHIR.
PaaS	*Platform as a Service* — a cloud operating model (relevant to hyperscaler lock-in, § 1.3 Dim. 2).
EMR / EHR	*Electronic Medical Record / Electronic Health Record* — the electronic patient record (within an institution / across institutions).
LIS	Laboratory information system.
PAS	*Patient Administration System* (chiefly NHS context).
FTE	*Full-Time Equivalent* — the unit of measure for staffing effort.
KPI	*Key Performance Indicator.*
RFI	*Request for Information* — market sounding ahead of a tender.
CDS	*Clinical Decision Support.*
PROMs	*Patient-Reported Outcome Measures.*

F-5 Organisations, Bodies and Standardisation Actors

ABBR.	EXPANSION
HL7	*Health Level Seven International* — the standardisation organisation behind FHIR, CDA and HL7 v2.
SDO	*Standards Development Organisation* — a body that develops and publishes standards. The term is used in this document for all such bodies, including openEHR International, IHE, SNOMED International and GS1. Where legal standing is what matters, the qualified term appears: a **formal** or **accredited** SDO develops its standards under a recognised accreditation and the balloting procedure attached to it — ISO and CEN by statute, HL7 International through ANSI accreditation. openEHR International and IHE publish through their own open procedures without such accreditation, and membership of the Joint Initiative Council does not confer it,

ABBR.	EXPANSION
	because the JIC coordinates rather than accredits (§ 1.5, points 2 and 5). The distinction is one of legal standing and not of technical quality; it bites where a regulation or a procurement rule refers to standards of a particular provenance.
CKM	*Clinical Knowledge Manager* — the platform on which openEHR archetypes are internationally reviewed and consented (§ 1.4, § 2.5, Appendix E-5).
OHDSI	*Observational Health Data Sciences and Informatics* — the community behind OMOP.
ADHA	*Australian Digital Health Agency* (Appendix E-4).
DHCW	*Digital Health and Care Wales* — the Welsh digital-health authority (Appendix D-7.2).
ÖGTelemed	Austrian Society for Telemedicine and eHealth — a **specialist society**, not a professional chamber (Appendix E-3).
FDMG	*Fachgruppe Datenmanagement im Gesundheitswesen* (Switzerland) — originator of the FHIR decision of 19 May 2026.
BAG	Federal Office of Public Health (Switzerland).
BMG	Federal Ministry of Health (Germany).
IKIM	Institute for Artificial Intelligence in Medicine, University Hospital Essen (Appendix D-7.4).
gematik	*Gesellschaft für Telematikanwendungen der Gesundheitskarte mbH* — the German national digital-health agency; specifies the ePA and, under § 373 SGB V, the hospital interfaces (Appendix D-10).
KBV	*Kassenärztliche Bundesvereinigung* — the National Association of Statutory Health Insurance Physicians; holder of the statutory mandate for the record-type specifications under § 355 SGB V (Appendix D-10.3).
mio42	mio42 GmbH — wholly owned subsidiary of the KBV, registered in October 2020, whose registered object is operational support for that mandate; it does not hold the mandate itself. The statutory mandate for the record types lies with the KBV (§ 355 SGB V).
Interop Council	*Interoperabilitätsrat* — the German interoperability council under § 385 SGB V with its expert and working groups; author of the core-profile position paper of March 2025 (Appendix D-10.3).
KIG	*Kompetenzzentrum für Interoperabilität im Gesundheitswesen* — the interoperability competence centre at gematik, which prepares the standards recommendations under § 385 SGB V.
INA	*Interoperabilitätsnavigator* — the national interoperability platform on which German standards, specifications and working-group papers are published.
RKI	*Robert Koch-Institut* — the German public-health institute; publisher of the notifiable-disease reporting profiles under § 14 IfSG (Appendix D-10.2).

F-6 Programmes, Initiatives and Infrastructures

ABBR.	EXPANSION
MII	Medical Informatics Initiative (Germany) — the alliance of university hospitals; established FHIR as the basis of the core dataset.
DIZ	Data integration centre — the MII unit at each university hospital that prepares care data for research.

ABBR.	EXPANSION
HiGHmed	An MII consortium; uses openEHR (ehrbase) as CDR backend.
NUM / CODEX	Network of University Medicine and its COVID data platform; a hybrid constellation of openEHR and FHIR as documented for 2020 to 2022 (Appendix C-3).
GECCO	*German Corona Consensus Dataset* — the COVID core dataset of German university medicine.
NDR	*National Data Resource* — the national health data platform of Wales; its Care Data Repository is a native FHIR store (Appendix D-7.2).
SwissHDS	Swiss Health Data Space.
MyHealth@EU	The EU infrastructure for cross-border exchange of health data.
EHDEN	*European Health Data & Evidence Network* — a European OMOP-based research network.
ELGA	The Austrian electronic health record — CDA-based.
ePA	The German electronic patient record; since ePA 3.0 of 15 January 2025 a hybrid of an IHE XDS.b document registry and FHIR data services (Appendix D-10).
ISiK	*Informationstechnische Systeme in Krankenhäusern* — the German hospital-interface specification of gematik, released in annual stages (Appendix D-10.2).
MIO	*Medizinische Informationsobjekte* — the KBV record types, published as FHIR packages `kbv.mio.*` (Appendix D-10.2).
eML	*elektronische Medikationsliste* — the electronic medication list of the ePA, fed automatically from e-prescription data (Appendix D-10.4).
DEMIS	*Deutsches Elektronisches Melde- und Informationssystem für den Infektionsschutz* — the German notifiable-disease reporting system (Appendix D-10.2).
TI	Telematics infrastructure (Germany).
FIKS	A major Norwegian openEHR-oriented EHR programme examined in a longitudinal case study (Christensen & Ellingsen 2016, Appendix D-5).
CatSalut / SISCAT	The Catalan health service and its integrated care system (Appendix D-2).
AUCDI	*Australian Clinical Data for Interoperability* — the Australian clinical dataset; per ADHA underpinned by FHIR **and** openEHR archetypes (Appendix E-4).
WHC	*Welsh Health Circular* — a binding circular of the Welsh Government (WHC/2023/018: the FHIR mandate).

F-7 Products and Platforms (Selection)

NAME	CLASSIFICATION
ehrbase	Open-source openEHR CDR (Apache 2.0); the backend of HiGHmed.
Better / Marand	A commercial openEHR platform (Slovenia); market leader in the openEHR segment.
DIPS Arena	A Norwegian HIS developed around an openEHR-oriented clinical model; the cited study documents the development programme rather than current market-wide deployment.

NAME	CLASSIFICATION
HAPI FHIR	An open-source FHIR server (Java).
Smile CDR / Aidbox / Firely Server	Commercial FHIR-native CDRs.
Firemetrics	A relational FHIR engine offered by a supplier with which an author of the Essen papers is affiliated (Appendix D-7.9).
CISTEC / KISIM	A Swiss HIS vendor and its product.
SAP IS-H / i.s.h.med	A widely deployed legacy HIS; its upcoming replacement is the occasion for architecture decisions in several countries.
ART-DECOR	A tooling environment for dataset and modelling governance in the HL7 environment. Serves its datasets and transactions as FHIR logical models through version-aware FHIR endpoints (Appendix E-7).

Note on Provenance

This living document states the current findings. A development log held by the author records source checks, changes between versions and statements retained with qualifications. Readers wishing to audit that history are referred to the log.

Character of the document. This assessment is evidence-based and analytical, and it ends in a procurement recommendation that is the author's own. It is not neutral between the two architectures for new infrastructure. Part I states the premise and limits of the recommendation; Part II presents the derivation, counterarguments and evidence boundaries. The decision remains with procurers and those politically responsible.

Method, and the use of a large language model. The literature search, the collation of sources, the drafting, and the consistency and readability checking across versions of this set were carried out with the assistance of a large language model. The editorial pass that set paragraph and sentence structure was made after the findings were fixed and did not change them. Not all of the cited literature was read in full: for part of it a summary extract was the working basis, and Appendix G of Document 4 records source by source which basis applies. Document 4 also states the limits of the method in full, including what a re-derivation of the evidence grades can and cannot establish. The grading rule, its application, the assessment and the conclusions are the author's, as is responsibility for any error that remains.

Transparency. The author is CEO of HL7 Deutschland, a Senior Expert and Community Event Manager at HL7 Europe, and a member of the ART-DECOR Expert Group. Those roles bear on a set of documents that recommends a FHIR-native default for new infrastructure, and they are stated for that reason.

The assessment presented here is his personal professional view as an interoperability expert and does not represent an official position of HL7 Deutschland or HL7 Europe. *The premise the recommendation rests on is stated in the open in Document 2 and in Document 3, so that it can be examined rather than inferred.*