Evidence Annex

Document 4 · Dr Kai U. Heitmann · Release 1.54.1 – 7 September 2026

Case material, comparison tables, source criticism, the source ledger and the bibliography.

Written for: Reviewers, researchers and readers auditing the source base.

Document 3 contains the argument supported by this annex. Documents 1 and 2 contain the decision and executive brief. Document 5 contains the German FHIR profile-corpus analysis, and Document 6 the desiderata addressed to the communities that could close the gaps recorded here.

References to numbered sections point to Document 3. Appendix D-10 is published separately as Document 5 and is therefore not reproduced here. Appendix F, the glossary, is included in Document 3.

How to use this annex. Appendices A to D document implementations, comparative evidence and production cases. Appendix E examines the character and interests of recent position and framework documents. Appendix G records the assessment basis, evidential role, grade, interests and retrieval status of each load-bearing source. The bibliography supplies the full reference record, and Appendix H registers the primary artefacts (legislation, specifications, manifests, queries) inspected directly.

Evidence grades apply to a proposition supported by a source, not to a publication as a whole. Deployment examples establish feasibility only to the extent stated. They do not by themselves establish effectiveness, superiority or lifecycle cost.

Consistency with Document 3. This annex supplies the evidence for the assessment and does not enlarge its conclusions. Five boundaries govern the reading of every appendix:

PROPOSITION CARRIED INTO DOCUMENT 3	WHAT THE EVIDENCE SUPPORTS	WHAT IT DOES NOT SUPPORT
FHIR has greater regulatory relevance for the European data exchanges examined	Current implementation and specification work uses FHIR profiles and implementation guides	That the EHDS Regulation prescribes FHIR as an internal persistence model
Persistent FHIR and openEHR repositories are feasible	Products and operational cases exist at institutional, regional or national scale	Comparative effectiveness, lower lifecycle cost or superior clinical outcomes
Each governed model conversion creates validation and maintenance obligations	Mapping studies, implementation reports and conformance findings document continuing work and possible semantic loss	A numerical relationship between the number of conversions and patient harm
Conformance does not ensure interoperability	Incompatible FHIR profiles and divergent declared openEHR product capabilities are documented	A commensurable score showing that one standard is more interoperable than the other
A FHIR persistence layer may remove one cross-standard boundary	The stored representation can also serve exchange where it is compatible with the required profile	That source-to-FHIR mappings or FHIR-to-FHIR profile transformations disappear

Architecture labels A–E follow § 1.6.1 of Document 3. They describe ideal types under stated assumptions, not exhaustive product categories. "FHIR-native" describes persistence only where the source says no more; it does not establish native primary capture.

Contents

Appendices

Appendix A: Repository Product Landscape (as of 2026)

Both standards are here described at the same layer — what products exist, on what licence, with what declared capability, so that the two pictures can be read against each other. The material behind them is not of equal quality, and the difference is stated in each part.

A-1: FHIR CDR Products

> **Core verdict.** FHIR-capable repository products exist across open-source, commercial and cloud-service delivery, and can be named.
>
> **Evidence grade.** D for the product entries — vendor documentation and product pages, without independent evaluation. The Aidbox performance figures below carry their own discounts and are read no further than they reach.
>
> **Primary limit.** Availability is not effectiveness. Nothing here establishes comparative outcomes, independent production maturity or clinical suitability, and no finding of the assessment rests on this list alone.

The following selection establishes that FHIR-capable repository products are available across open-source, commercial and cloud-service models. It does not establish comparative effectiveness, independent production maturity or suitability for a particular procurement. Production cases are assessed separately in Appendix D-7.

PRODUCT	DELIVERY MODEL	RELEVANT REPOSITORY CAPABILITIES
HAPI FHIR JPA Server	Open source, Java, Apache 2.0	Relational persistence and a widely used reference implementation
Firely Server	Commercial, .NET	FHIR persistence with MongoDB or SQL Server backends
Smile CDR	Commercial, HAPI-based	Clustering, subscriptions, master-data functions and SMART on FHIR hosting
Microsoft Azure Health Data Services / FHIR Service	Managed cloud service	FHIR persistence, SMART on FHIR and Bulk Data Export
Google Cloud Healthcare API	Managed cloud service	FHIR R4/R5 services and BigQuery export
Aidbox	Commercial	FHIR repository with GraphQL and SQL-oriented access
Medplum	Open source, Apache 2.0	FHIR repository and application-development components
LinuxForHealth FHIR Server	Open source, Apache 2.0	Server implementation originating from IBM FHIR Server

Devitt's July 2026 catalogue lists roughly thirty FHIR server products. It is useful as a market inventory but remains an informal, volatile source from a single author and is graded D.

What comparison material exists, and what class it belongs to. Two comparisons of FHIR repository products were assessed for this appendix, and neither is peer-reviewed.

- **An academic-network evaluation, and now dated.** The CTSA data-harmonization project published a comparison of seven products — Bunsen, the Microsoft FHIR Server open-source and managed editions, the Google Cloud Healthcare API, Vonk, HAPI FHIR and Smile CDR — across seven criteria: REST APIs, authentication, bulk-insert endpoints, FHIR version, control over the database and its location, security and compliance, and external terminology server. The evaluators state their own boundary plainly: they set each server up from the available documentation without expert consultation. It reaches no overall recommendation and records instead where each product fell short, including identifier masking in two managed services and absent R4 support in several. Its reference points are HAPI 3.8.0 and 4.0.0, which places it in 2019. **Evidence grade D**, and stale, carried here as a record that the exercise was attempted in a research network, not as a current picture.
- **A vendor benchmark, with an open harness.** Health Samurai published on 29 June 2026 a throughput comparison of Aidbox, HAPI FHIR, Medplum and the Microsoft FHIR Server on one 64-core machine with local NVMe storage, each application container given 8 vCPU, against a Synthea dataset of 1,000 patients and roughly two million R4 resources, driven by Grafana k6. Reported: CRUD throughput of 5,212 requests per second for Aidbox, 3,058 for HAPI, 1,420 for Medplum and 440 for the Microsoft server; search throughput 3,404, 1,005, 1,796 and 261 respectively; batch import 2,678, 2,214, 764 and 448 resources per second; and database size for the same content ranging from 4.24 GB for the Microsoft server to 22.6 GB for HAPI. **Three discounts and one credit.** The publisher builds Aidbox, Aidbox leads on every axis but storage, and the authors say themselves that they know their own product best. The dataset is 1,000 patients, which is not the scale the title claims and is well below any national deployment. And a single-machine, single-configuration run measures the configuration as much as the product. The credit is real and unusual: the harness, the code and the results are published on GitHub and the report is regenerated daily, so a third party can rerun it and contest the numbers. **Evidence grade D** for the ranking, which a vendor produced about itself; the open harness is what makes it worth listing rather than raising the grade.

An asymmetry, and it runs against expectation. The openEHR half of this appendix rests on a peer-reviewed survey of seventeen identified products with a published answer set. The FHIR half rests on vendor documentation, a seven-year-old evaluation and a supplier benchmarking itself. **No peer-reviewed comparative survey of FHIR repository products was found in the material assessed for this document.**

That is a statement about this search rather than about the world, and it may simply be that such a study exists and was not reached. But the search was not casual. The shape of the result is the opposite of what the relative sizes of the two ecosystems would suggest: the smaller community has the better-documented product layer.

The consequence for this appendix is that **A-1 is the weaker half and its findings carry less than A-2's**, which is stated here rather than left for a reader to infer. The consequence for the field is Desideratum 22.

The application layer, for contrast, has been surveyed. Griffin et al. (2022) collected 112 real-world FHIR applications from public repositories and a developer survey. The breakdown is 74 for clinical care and 45 for research, with 56 implemented across more than one site. Also recorded are 55 using SMART on FHIR, 69 on R4, 67 stand-alone web applications and 51 embedded in an electronic health record.

The figure worth carrying is that **49 of the 112 appear in no electronic-health-record application gallery**. That says something about how much of the ecosystem any catalogue can see, including the catalogues above. It is a brief communication with a convenience sample and self-reported data, several authors are officers of HL7, and it describes applications rather than repositories. **Evidence grade C**, used for breadth and for that one figure.

A-2: openEHR CDR Products — a Peer-Reviewed Survey (Delussu et al. 2024)

What it is. A survey article in the *International Journal of Medical Informatics*. The authors searched the openEHR vendor listings, the Discourse forum, the web and personal contacts for openEHR-based repositories mature enough for production use in hospitals or research settings. They identified **seventeen**, and put a fifty-four-question instrument to their vendors and developers. **Eleven answered.**

The questions cover licensing, installation, ecosystem tooling, implementation detail, operations on the openEHR elements, adherence to the specifications and sharing capability. The collected answers and the analysis code are published for reuse. It is the successor to a 2013 catalogue by other authors, which allows a ten-year comparison on five fields.

Source character. Peer-reviewed, method disclosed, raw data and scripts open: the strongest source in this document for the openEHR product layer.

Three things bound it. Every cell of the two result tables is a **vendor self-declaration** that the authors did not verify. **Thirty-five per cent of the suppliers approached did not respond or did not complete the instrument**, which the authors record as their first limitation.

And the team builds openEHR tooling and is openEHR-aligned, which makes the findings running against openEHR's own claims stronger and the favourable ones weaker. Data current to 2023. **Evidence grade B** for the declared product properties; **grade C** for any inference from a declared property to what a deployment can do.

What it establishes in openEHR's favour. The count of production-capable openEHR repositories rose from eleven to seventeen in the decade since the 2013 catalogue. Supplier countries new to the field contribute five products from China and one each from India and Mexico. AQL has overtaken SQL as the query language, reversing the 2013 ratio of eight to five. Docker images are available for ten products and cloud-managed instances for eight.

Named deployments include the National Digital Service in Scotland and a German COVID-19 data platform (EHRBase), and Moscow city, Slovenia and Malta (Better). Others are twenty deployments in the Netherlands (Base24), five of seven Norwegian university hospitals (ArenaEHR) and a thirty-hospital installation (EHRDB). The United States is absent from the list altogether, which the authors attribute to a market held by vendors on closed standards.

What it establishes about portability, which is the part this assessment carries. The case for openEHR as the vendor-neutral option rests on two mechanisms that belong to the specification rather than to any product. The first is the Archetype Query Language, which is meant to make a query portable across implementations. The second is the EHR Extract, the specified route for taking a record out of one openEHR system. Across the eleven products that answered:

PROPERTY AS DECLARED BY THE VENDOR	PRODUCTS, OF 11
Query language AQL	6
Query language SQL	4
Neither — a query formalism of the vendor's own	2
REST API mostly conformant to the openEHR specification	4, plus 1 in development
openEHR EHR Extract supported	1, plus 1 through a separate product
Free to use and open source	2
ADL 2.0 supported	1
HL7 FHIR import or export in production	3, plus 3 in development or through a separate integration engine

One product declares both AQL and SQL, which is why the query rows sum to twelve.

Serialisation is dispersed in the same way. Compositions move as canonical XML, canonical JSON, FLAT, STRUCTURED, a vendor's Template Data Document or a vendor's AQL-path/value JSON. Templates move as an Operational Template or as a Simplified Data Template. Reference-model versions run from 1.0.1 to 1.1.0 with no common floor. ADL stands at 1.4 for nine products, 1.5 for one and 2.0 for one, a decade after ADL 2 was specified.

What the authors conclude from their own dispersion figures. They write that the differences between the flat formats implemented in the various repositories "prevent the straightforward data interoperability between the various CDRs" unless an external intervention is applied. They attribute this in part to the simplified data-format specification still being in development.

They then name what is missing to close it: a "testing framework to assess the compliance of openEHR systems and products" to the specifications. They hold such a framework essential to the success of the standard. The openEHR conformance guide is itself still in development.

Where this is used, and where it is not. It is used at § 2.6.2 to put measurements behind two objections the assessment previously weighed without them. It is used at § 2.6.1 to bound the reach of the AQL argument to the products that implement AQL, and at § 1.5 as the openEHR-side counterpart to the German conformance finding. It also supplies product context for the openEHR deployments graded in Appendix D-3.

It is not used to argue that openEHR is less interoperable than FHIR. The two findings are not commensurable. The German corpus counts artefacts across four specification families, this survey counts declared properties of eleven products, and neither measures what a deployment does in the field. What they share is the shape of the result, and that is what is carried forward. Conformance to an open specification is not by itself interoperability, on either side of the debate.

Appendix B: Detailed Observations from the openEHR Tutorial (Pazos et al. 2015)

What this source is, before anything is taken from it. A tutorial presentation given inside the openEHR community in 2015. It is teaching and advocacy material, not a controlled comparison: it names no protocol, no competing configuration and no repetition, and its figures were produced by parties with an interest in the subject. Its technical horizon is a decade old — the products, storage engines, indexing strategies and query implementations it describes have all moved since. **Evidence grade C** for the statements of position drawn on at § 2.2; **D** for the storage and query figures.

What it is used for, and the one thing it must not be used for. It is used for what the openEHR community said about its own architecture at the time. Persistence is deliberately left to the implementer, no single solution fits all cases, and hybrid arrangements were already being recommended from inside the camp. Those are statements of position, and a presentation is a sound record of a position.

The storage and query figures below are not transferable to any product in use today, and no proposition in this assessment rests on them as a current performance claim. Where current repository performance figures appear, they are at Appendix A-1, they are from 2026, and they carry their own discounts.

The following slides underpin the architecture statements of § 2.2:

- **Slide 16** — "The choice depends on the context": the choice of database technology rests with the implementer.
- **Slide 21** — performance and size tests on 4.2 million patient records: openEHR XML storage was reported to require markedly more storage space than an epidemiology-optimised relational reference schema, by a factor of roughly 10 to 40, and population queries were slower without specific indexing. **As a 2015 figure from a**

tutorial, this records what was measured then on that configuration and says nothing about any current product.

- **Slide 37** — "openEHR doesn't define how to store data": persistence is explicitly left to the implementation.
- **Slide 50 ff.** — AQL examples: a powerful declarative language, productive use predominantly via specialised editor tools.
- **Slide 76** — conclusions: "There is no one-fits-all solution"; hybrid architectures are recommended.

Appendix C: Observations from the Swiss and German Hospital Landscape (as of August 2026)

In Switzerland a split picture is emerging. Several recent procurement decisions by large hospitals have gone to integrated commercial HIS platforms, most visibly Epic, while individual university hospitals and the HIS market leader have taken up openEHR components. **No market share is claimed here.** The cases below are those that were publicly documented and found. Unless stated otherwise, the entries below rest on press reports and institutional communication and carry evidence grade D. There is no procurement dataset behind them, no defined denominator of "large hospitals" and no fixed observation window. The pattern is therefore a set of named decisions rather than a measured distribution.

C-1: Where openEHR Is Actively Deployed or Being Expanded

- **University Hospital Basel (USB):** In early 2025 commissioned a consortium of OWT (a Swisscom subsidiary), Swisscom, x-tention and Better to build an openEHR-conformant data platform; explicit separation of the application layer and the data platform, for vendor independence of the medical data.
- **CISTEC (KISIM/KISIM+):** The Swiss HIS market leader describes its approach as "KISIM on Top of openEHR". Important for classification: not a complete openEHR re-architecture, but a hybrid strategy in which openEHR is integrated alongside FHIR/HL7 v2 as one of several standards (concrete modules: PROMs, reports, medication, the Switzerland-wide diagnosis list).
- **Kantonsspital Aarau:** According to industry reports, openEHR primarily for secondary use (a data lake for research, studies, AI training), not as a HIS replacement.
- **openEHR Switzerland:** Founded in March 2024 as a national association; membership includes CISTEC, Better and hospital representatives.
- **openEHR Switzerland and HL7 Switzerland (2026):** A joint working group is developing a reusable implementation-guide "blueprint" from a vaccination showcase of the Swiss GovTech Hackathon, pairing openEHR persistence with CH-VACD/FHIR exchange via the FHIRConnect/openFHIR transformation. An announced proof of concept, not a productive deployment (see Appendix E-8; evidence grade D). It explores the internal-persistence side that the FDMG decision deliberately left open.
- **Federal Quality Commission (EQK):** First productive openEHR project for the standardised capture of patient feedback (PROMs).

C-2: Where Hospitals Chose Epic Without a Documented openEHR Evaluation

The Insel Gruppe introduced Epic in March 2024 after a four-year project phase. In February 2026 it disclosed total costs of around CHF 228 million: CHF 182.5 million introduction, CHF 45 million operation until 2032. The system replaced around 50 predecessor systems. No public source found in this search indicates that openEHR was formally evaluated as an equivalent architectural option, and the tender's focus was on a turnkey overall solution.

That is a statement about what is public and what was found, not a finding that no evaluation took place. Procurement deliberations are not generally published, and their absence from the record cannot be read as their absence in fact. The Insel decision is therefore not a documented case of an openEHR rejection after evaluation, but an example of a tendering logic in which openEHR stands structurally outside the assessment frame.

Similar paths are documented for CHUV Lausanne, USZ Zürich and the Lucerne cantonal hospital. In August 2025 it became known that the Canton of Bern intends to equip all listed hospitals with Epic. Noteworthy is the contrast with USB Basel. There, in 2023, CISTEC (KISIM) and Epic were evaluated against each other, two complete HIS products, not openEHR as an option in its own right.

C-3: German Cases as a Comparative Perspective

The first item below is a procurement case. The second is not: a publicly funded research data platform is not a tender, and it is set out here because it is invoked in the debate as though it were a decision between architectures. What the sources carry, and where they stop, is therefore stated at some length.

- **Charité Berlin → Epic (December 2025):** Europe's largest university hospital concluded the award to Epic after a roughly two-year procedure (around €200 million over ten years), following the announcement of SAP's withdrawal from the HIS market. It is the first productive Epic installation in a German hospital. openEHR does not appear in the public award documents as an assessment option in its own right; the competition ran between complete HIS products.
- **NUM CODEX and HiGHmed:** HiGHmed is a consortium of the German Medical Informatics Initiative whose architecture includes openEHR. NUM CODEX was established in 2020 as a national COVID-19 research data platform and reused parts of the technological and organisational environment developed within the Medical Informatics Initiative and HiGHmed. The two initiatives are related but are not identical, and the composite wording "NUM CODEX / HiGHmed" current in the debate obscures that.

What is documented, and for when. For the period 2020 to 2022 the GECCO dataset was represented in both HL7 FHIR R4 and openEHR, and both representations were treated as applicable to the COVID-19 research use case (Behrend & Ingenerf, GMDS 2022). A conversion service enabled data held in openEHR repositories to be transformed into the GECCO FHIR representation required for exchange (Ladas et al. 2022).

The evidence supports the parallel use of the two standards. FHIR served as the common exchange and secondary-use representation, openEHR as a local persistence and modelling approach at the participating sites that had adopted it. Evidence grade B.

What the sources do not establish. The current operational status of the openEHR components within NUM CODEX is not established. The 2023 architecture description of the NUM research data platform (Heyder et al. 2023) sets out the GECCO dataset, its FHIR specification and federated FHIR repositories at the data integration centres. It does not describe openEHR repositories, ehrbase or a continuing conversion bridge.

The openEHR-side components of the NUM code base were last updated between May 2024 and February 2025. The repository carrying the broker function was archived on 8 April 2024 with the note that it is no longer maintained as an official ehrbase component, the work having moved to a NUM fork itself last touched in October 2024.

Repository activity and archival notices indicate changes in maintenance responsibility and establish neither continued production use nor decommissioning. Evidence grade C for the repository observations, which are self-reported project metadata rather than a statement of operational status. No conclusion about the operational role of openEHR in 2026 therefore follows from the available public evidence.

The sequence of FHIR adoption in the Medical Informatics Initiative. FHIR use in the initiative predates GECCO, and three dated steps lie behind it. The first version of the MII core dataset was adopted in March 2017 as a textual description. By 7 December 2018 the FHIR work on the base modules was under way: a cross-consortium mappathon on 9 November 2018 had mapped example care data into HL7 FHIR, and FHIR profiles for Encounter, Condition, Procedure, Observation, DiagnosticReport and MedicationStatement were named as the deliverables of the first phase of the core-dataset development, as the presentation of the initiative's interoperability working group at the HL7 Deutschland Interoperabilitätsforum of that date records (Ammon 2018; `material/HL7DE-MII-`

`Kerndatensatz-2018-12-07.pdf`, slides 10 and 11). In April 2019 the initiative's National Steering Committee decided formally that HL7 FHIR is the standard for the technical representation of the core dataset, and at the same time that ART-DECOR would serve the dataset modelling and Forge and Simplifier.net the authoring and publication of the profiles (MII, *Manteldokument* of the core-dataset implementation guide, Simplifier.net; the initiative's own account, evidence grade B). An earlier version of this passage placed the FHIR decision in March 2017 together with the dataset itself; the two are separate decisions two years apart, and the profiling work of 2018 preceded the formal decision rather than following it.

GECCO, developed in 2020, drew on those existing MII FHIR profiles, as its own publication records (Sass et al. 2020). The present author is a co-author of that paper, and the statement cited is the paper's own record of its inputs.

GECCO therefore reflects the prior adoption of FHIR within the MII rather than explaining it. The sequence runs the opposite way to the claim, current in the debate, that FHIR predominates in the core dataset because GECCO was specified in FHIR.

What the case supports, and what it does not. The case provides evidence of a function-oriented coexistence of standards during the documented period: FHIR for common exchange and research-data integration, openEHR for detailed local modelling and persistence where it had already been deployed.

It does not provide evidence of a formal migration from openEHR to FHIR, of an abandonment of openEHR, or of the present operational status of the openEHR components. The reading sometimes offered that the NUM "pivoted from ehrbase/openEHR to the MII FHIR architecture" rests on no documented decision, and no announcement of such a switch was found.

C-4: Recurring Arguments Against Choosing openEHR as the HIS Core

From industry reports, and not from any single methodically documented study, five recurring arguments crystallise (to be read as interpretative industry observation, with the exception of point 2, which is empirically supported by Christensen/Ellingsen 2016):

1. **No turnkey HIS.** openEHR supplies data architecture and archetypes, but no finished user interfaces, no DRG billing logic, no bed planning. Large institutions avoid the integration responsibility this entails.
2. **Learning curve and skills shortage.** Two-level modelling requires specialised staff; the Swiss market for it is small. In the FIKS project, only 60 active users were registered 18 months after the CKM launch.
3. **Maturity of the commercial ecosystem.** The openEHR platform market is characterised by smaller vendors; the vendor-survival risk is assessed negatively for investments in the hundreds of millions. (Qualified by CISTEC's entry and the USB consortium after 2024.)
4. **Performance concerns for real-time queries.** The deep, nested structure requires considerable indexing work in real-time scenarios.
5. **Prioritisation of data exchange.** Many hospitals argue that a conformant external FHIR interface suffices, which corresponds to the FDMG decision.

C-5: Overall Picture

A direct standard-to-standard evaluation "openEHR versus proprietary HIS data model" followed by a reasoned rejection is documented in neither Swiss nor German sources. openEHR mostly falls through the requirements grid of monolithic HIS tenders without ever competing there as an equivalent option. The statement "openEHR was evaluated and rejected" should therefore be used with caution; it is more precise to say that openEHR stands structurally outside the assessment in procurement logics aimed at complete systems.

Appendix D: Comparison Tables, Production Cases and Studies

D-1: Consolidated Standards Comparison after Oemig/Blobel

ASPECT	HL7 V2.X	HL7 V3 / CDA	HL7 FHIR	OPENEHR	OMOP
Publisher	HL7 Int.	HL7 Int.	HL7 Int.	openEHR Foundation	OHDSI
First release	1987	2001/2005	2014 (DSTU1)	ca. 2000	2010
Current maintenance	maintained (v2.9.1)	discontinued since 2018	active (R4, R5; R6 in progress)	active	active (CDM v5)
Adoption (2018–2020)	>1000	declining, >100	growing strongly, >1000	small	small, primarily research
Data-model level	Physical	Logical	Logical (computational)	Logical (information)	Physical
Persistence prescription	none	none	not prescribed, increasingly native in CDRs	explicitly none	fixed relational structure
Query language	implicit	OCL	FHIR Search, FHIRPath	AQL	SQL against the CDM schema
Tooling (open source)	HAPI v2	ART-DECOR, MDHT	HAPI FHIR, Firely, SUSHI, Inferno, ART-DECOR	CKM, Better Studio, ehrbase	OHDSI tools
Suitability as persistence layer	no	no	yes (native FHIR CDRs)	yes (primary purpose)	yes (research-specific)

A note on the tooling row. In the source table ART-DECOR appears only under HL7 v3/CDA. That classification does not hold: the environment also serves ValueSets, CodeSystems, ConceptMaps, NamingSystems, Questionnaires and logical models through version-aware FHIR endpoints (Appendix E-7), and the entry above is extended accordingly. The point is minor in itself and illustrates the weighting problem named below: a tooling row compiled at one point in time ages, and ages asymmetrically.

Oemig/Blobel (2020) conclude that, on an objective assessment of all aspects, FHIR is to be preferred. This assessment does not adopt that conclusion unexamined, because the criteria are heavily weighted towards those properties in which FHIR structurally leads (adoption, tooling, web-nativeness). The analytical value of the source lies in the structured classification by data-model levels.

D-2: Catalonia/CatSalut — Chronology and Source Overview

PERIOD	EVENT	SOURCE TYPE
2018–2022	SISCAT master plan establishes openEHR as strategic basis	Official CatSalut documents
July 2021	RFI for platform components	CatSalut/CTTI, public
Sept. 2022	RFI final report (15 offers)	openEHR Discourse, public

PERIOD	EVENT	SOURCE TYPE
Dec. 2022	Tender for clinical knowledge-management platform	CatSalut/CTTI, public
June 2023	CTTI award: IBM / vitagroup (ehrbase/HIP CDR) / Red Hat	Press releases IBM, vitagroup, openEHR Foundation
June 2023	Award T-Systems/Better for EPMA (60+ hospitals)	Better, T-Systems, digitalhealthnews.eu
2023–2026	Rollout/operation (scope not independently verified)	Vendor PR, institutional communication
2025	Book chapter Piera-Jiménez & Carot-Sans, Springer	Edited academic volume (Springer); authors = CatSalut officials
Apr. 2025	Tender for platform services (€7.1 million)	openEHR Discourse, public
Oct. 2025	Partnership Tietoevry Care / NTT DATA	Press release

Evidence gap: In the material identified, the reports are procurement or planning documents, vendor publications or institutional self-presentations; no independent outcome evaluation was found. The first author of the assessed Springer chapter (Piera-Jiménez, 2025) is now CEO of openEHR International, so the chapter is treated as institutional self-description within an edited academic publication rather than as independent outcome evidence (§ 1.4.1).

D-3: European openEHR Reference Implementations by Evidence Quality

REGION / PROJECT	STATUS (CA. 2026)	PEER-REVIEWED OUTCOME DATA	CHARACTERISATION OF SOURCES
Norway FIKS (Christensen & Ellingsen 2016)	Completed / documented	Yes — Int J Med Inf 2016	Independent four-year longitudinal study; critical findings; methodologically strongest source
Slovenia (Better)	Ongoing	No	Vendor PR
Catalonia (CatSalut)	Rollout since 2023	No	Vendor PR, institutional self-presentation, Springer chapter by the officials
Finland (parts of THL)	Partial / ongoing	No	Project reports
Sweden (Cambio)	Ongoing	No	Vendor communication
USB Basel (Switzerland)	Build-up since 2025	No	Vendor PR
NUM CODEX/CRR (DE)	Documented for 2020–2022; present operational state not established	Partial — Heyder et al. 2023, which does not describe the openEHR side	Hybrid FHIR+openEHR constellation documented for 2020–2022 (Appendix C-3); no evidence of openEHR-specific interoperability gains

Conclusion: The search identified one peer-reviewed longitudinal case study that directly examines the implementation logic of a large-scale openEHR programme. It presents a differentiated and partly critical picture.

The other cases assembled here do not provide comparable independent longitudinal evaluations. This is a finding about the reviewed evidence base, not proof that no other evaluation exists.

D-4: European FHIR Reference Implementations by Evidence Quality

Preliminary remark: why the comparison is not fully symmetric.

Two countervailing reasons must be made transparent.

First (qualifying the FHIR advantage): FHIR is required or selected by several national and European implementation regimes, including specifications under § 355 SGB V and current EHDS/MyHealth@EU implementation work. The precise legal mechanism differs between jurisdictions: the EHDS Regulation does not itself prescribe an internal FHIR persistence model. Regulatory adoption demonstrates policy and implementation relevance, not superiority or effectiveness.

Second (qualifying the openEHR disadvantage): The smaller number of documented openEHR deployments is partly structural: the installations are fewer and smaller. That creates an asymmetry of scale and maturity, not necessarily of quality; the one peer-reviewed longitudinal study identified is on the openEHR side (D-5).

REGION / PROJECT	STATUS (CA. 2026)	PEER-REVIEWED OUTCOME DATA	CHARACTERISATION OF SOURCES
MyHealth@EU / eHDSI	Operational since 2019; growing	Partial — Bruthans & Jiráková 2023	KPI analysis of public EU transaction data; descriptive; no causal benefit assessment
Germany — ePA / TI	Opt-out rollout from Jan. 2025	Partial — adoption studies 2024/2025	Usage statistics; technical interoperability gains not independently evaluated
Germany — e-prescription	Mandatory since Jan. 2024	No	gematik/BMG communication
Germany — MII core dataset	Operational in ≥36 university hospitals	Partial — Ammon et al. 2024; Heyder et al. 2023 (NUM platform)	Working-group and architecture reports; no FHIR-specific interoperability evidence (D-7.4)
Finland — Kanta (FHIR transition)	Migration ongoing	Partial — Hyppönen et al. 2019	Longitudinal usage data (still CDA-based)
NHS England — GP Connect	Productive since ca. 2019–2021	No	Governance documents; no independent outcome study
Wales — National Data Resource / Care Data Repository	Mandate since June 2023; CDR productive, being built out	No	Government circular (WHC/2023/018), DHCW implementation guide, HL7 UK article; native FHIR store; no independent evaluation
Switzerland — SwissHDS	Decision May 2026	No	FDMG decision, EY report; no operational system
Austria — ELGA (FHIR migration)	Early phase	No	ELGA reports
Spain/Madrid — INFOBANCO	Ongoing	Partial — Pedrera-Jiménez et al. 2023	Technical feasibility demonstration; no long-term outcome data

Conclusion. For FHIR too, independent longitudinal studies with outcome data are largely absent from the reviewed material; most sources are government communication, agency reports or vendor publications. The KPI analysis of eHDSI identified here (Bruthans & Jiráková 2023) is descriptive and notes that the programme had predominantly test character during the observation period. The clearer difference is regulatory and implementation adoption, not demonstrated effectiveness.

D-5: Empirical Case Study — Christensen & Ellingsen (2016) on the FIKS Project

The peer-reviewed study *"Evaluating Model-Driven Development for large-scale EHRs through the openEHR approach"* (Int J Med Inf 89, 2016, pp. 43–54) is the methodologically most careful published long-term case study of a large-scale openEHR implementation.

Study design. Four years of interpretative field research (2012–2015) on the FIKS project of the Northern Norway Regional Health Authority, with DIPS as the main vendor: 440 hours of participant observation, 22 interviews and extensive document analysis. The paper reports that DIPS then held 82% of the Norwegian hospital EHR market. That figure describes the vendor's market position at the time, not the deployment share of DIPS Arena or of an openEHR-based product. The project budget approached €90 million.

Assumptions examined and findings.

1. *Clean separation of clinical and technical concerns* — **refuted.** Both activities required both kinds of competence.
2. *Clinical self-governance of modelling* — **not realised at large scale.** Responsibility migrated from the local via the national to the international level; local users in the end had "virtually no control".
3. *Broad user commitment* — **markedly undershot.** Only 60 active users 18 months after the CKM launch; recruiting proved difficult.
4. *Modelling as a simple mapping of existing practice* — **strongly qualified.** Modelling was a complex standardisation process that itself changed clinical practice.

The planned surgery-planning module (October 2012) was still not available in November 2014, primarily for lack of approved archetypes.

Significance for this analysis. The study provides strong evidence that large-scale model-driven development created substantial modelling, coordination and recruitment challenges in the case examined. It does not measure the cost of an openEHR–FHIR bridge and should not be cited as doing so.

The failure of the four assumptions is not evidence that openEHR itself failed. The authors argue for more formal modelling structures rather than abandonment. *Methodological limitation:* a single case in a specific national setting, stronger representation of the vendor perspective and no coverage of the project after 2016.

Evidence grade A for the case findings, with the qualification that it examines one project rather than comparative architectural fitness.

D-6: External Industry Analysis — Gartner Research Note (2017)

The Research Note *"Evaluate openEHR Standards for Managing Clinical Content Across the Care Continuum"* (April 2017) is the most important external industry-analytical counterpoint to the FHIR-leaning materials.

Sourcing boundary. The report exists and is cited independently of this assessment, by title and date, in a European Commission staff working document (SWD(2018) 126 final of 25 April 2018, footnote 51). It is drawn on there for the proposition that proprietary solutions and vendor lock-in obstruct interoperable cross-border records. No currently accessible full-text copy could be obtained for this assessment.

What follows is therefore reproduced from the surviving summary of the report and from clearly identified secondary descriptions of it, not from the full text. The wording of the recommendations, scenarios and cautions

below is summary wording rather than quotation. The author attribution (Jones and Hallenberg) and the document number G00317087 come from that summary and are not corroborated by the Commission citation, which gives title, month and publisher only.

Until the complete report has been obtained and read, nothing beyond the summary should be attributed to it. No finding of this assessment rests on this source alone: it is carried as a counterpoint, and each proposition drawn from it is stated as the analyst position it is. **Evidence grade C/D**: an analyst product without peer review, read in summary only.

Recommendations for CIOs. Reduce vendor lock-in by persisting structured clinical content outside the proprietary EHR; evaluate openEHR as a means towards a "universally consistent and semantically interoperable clinical data model" in the target architecture, especially for national/regional e-health initiatives; invest in clinical informatics skills.

Clarifications (Gartner). openEHR is not to be confused with an open-source EHR product, a complete megasuite alternative, or a messaging standard. It can be operated centralised or federated.

Three suitable scenarios (all presupposing mature governance and a long-term vision): a national e-health programme with a 5–10-year horizon; an HDO federation under value-based care; a municipality with a fragmented multi-EHR landscape.

Caution (Gartner). openEHR is not suitable for organisations that have not yet clarified their long-term ecosystem vision or cannot invest in clinical informatics skills. Platform vendors must be examined for vision, ability to execute, genuine openness and ISO conformance; contractual data sovereignty at the end of the contract must be regulated.

Significance for this analysis. The source validates the pro-openEHR arguments of § 1.4 as the position of a prominent independent industry analysis, but sharply qualifies them: openEHR is not suitable for every organisation. The documented procurement cases in which openEHR plays no role at the HIS core (Inselspital, Charité, USZ, CHUV, Lucerne) fit the Gartner logic: these institutions aim at monolithic megasuites with operational stability, not at building an open ecosystem.

D-7: FHIR-Native CDRs in Production — State of the Evidence

> **Core verdict.** FHIR-native persistence runs in production at national and hospital scale. Wales is the clearest documented national case; the German vendor market is weaker than its labels suggest.
>
> **Evidence grade.** Mixed, B down to D. Programme documentation and published descriptions where they exist, vendor self-declaration where they do not; each entry carries its own grade.
>
> **Primary limit.** No independent, peer-reviewed effectiveness study was identified, the same limit that applies to openEHR at Appendix D-3. Feasibility is documented; comparative effectiveness is not. Two of the products examined describe themselves as FHIR-native and place themselves at Option B (D-7.7, D-7.9).

This appendix documents the production instances that underpin the claim that FHIR-native CDRs are production-ready (§ 2.4; Part I, Finding 3) — with expressly limited evidence quality.

Preliminary remark on symmetry. For FHIR-native persistence, the same evidence critique applies as for openEHR (Appendix D-3): no independent, peer-reviewed effectiveness study was identified in the reviewed material. The cases demonstrate feasibility and operation to the extent stated, not superiority.

D-7.1 Overview

A note on the term. "FHIR-native" in this appendix describes **where the data are kept**, not where they are recorded. Every case below is fed from source systems using another model, so in the five-way scheme of § 1.6.1 these are instances of **Option B**, not Option C.

CASE	LEVEL	CHARACTER	EVIDENCE GRADE
Wales — NDR / Care Data Repository	National	Persistent, native FHIR store as central data platform	C (government document + trade article; no independent evaluation)
MII / data integration centres (DE)	Infrastructure, 36+ sites	FHIR as implementation basis of the core dataset	B (peer-reviewed, but institutional)
UH Essen — DermaDashboard	Institution	Oncological cohort building on a relational FHIR store	C (peer-reviewed, but a viewpoint)
UH Essen — blood-product dashboard	Institution	Clinical-operational real-time workflow on a FHIR store	C (peer-reviewed short report)
MII / FDPG national aggregate	Infrastructure, 41 sites	Pooled FHIR-structured holdings (3+ bn lab values; 27+ M persons)	C (self-reported public aggregate, not independently audited)
DIZ Essen — infrastructure size	Institution	FHIR store in the billion-resource range	D (self-reported, not independently verified)
Turkey — national repository	National	Reported ~38 billion FHIR resources on ~32 M patients	D (conference presentation by the implementing party)
Ukraine — national record	National	FHIR-native persistence, non-FHIR REST interface; ~39 M patients	D (personal communication, HL7 UK; no published primary source)
Lombardy (IT) — regional CDR	Regional	FHIR-based clinical data repository, ~10 M patients	D (conference presentation by the implementing party)
Singapore — Healthier SG	National programme	FHIR implementation at national programme scale; persistence model not determinable	D (conference presentation by the implementing party)

D-7.2 Wales — National Data Resource / Care Data Repository

Regulatory frame. Welsh Health Circular WHC/2023/018 (14 June 2023) obliges all NHS Wales bodies to adopt HL7 FHIR as the foundational standard for interoperability. The Welsh FHIR profiles are maintained as the *Wales FHIR Implementation Guide* (currently v2.6.0) and are derived from the UK Core.

Architecture. The National Data Resource (NDR) replaces fragmented local databases with a unified national platform. Its structural centrepiece, the Care Data Repository (CDR), processes FHIR messages and is implemented as a **persistent FHIR store**. According to Digital Health and Care Wales, FHIR there is explicitly not a temporary integration or translation layer but the long-term structure of the data itself; every clinical and administrative record in the CDR exists as a native FHIR resource.

Productive content. Entry was made via electronic prescribing in the secondary sector, whose success provided the basis for the national FHIR mandate. Now implemented are, among others, administrative data (demographics, organisation, location), structured and coded profiles for drug allergies, and paediatric growth charts.

Relevance to the architecture analysis. Wales is a documented national-scale case of a persistent FHIR CDR. It demonstrates that Option B (§ 1.6.1) can be implemented across a health system without an openEHR persistence layer. It is not Option C because care is recorded in source systems and the repository is populated from them.

Limits of the evidence.

- No independent, peer-reviewed evaluation; the account stems from an HL7 UK trade article with a DHCW official as interviewee and from government and agency documents. The source is thus institutional self-presentation, not external assessment. **Evidence grade C** for the production case; A for what the circular requires; B for the implementation guide.
- The mapping has not vanished: data from the primary systems (including PAS) continue to be delivered via HL7 v2 and transformed into national FHIR profiles. The FHIR-native CDR eliminates the *second* transformation stage (persistence → exchange), not the *first* (source → target model).
- The implementation status is incremental; a complete longitudinal patient record is the goal, not the achieved state.
- No documented comparison with an openEHR-based alternative architecture, and Wales, as far as publicly visible, did not formally evaluate openEHR as an equivalent option. The same qualification therefore applies here that Appendix C-5 formulates for the Swiss Epic awards, with the sign reversed.

D-7.3 University Hospital Essen — Two Clinical Use Cases

An interest disclosed at D-7.9 attaches to both papers below: a co-author of each is affiliated with a supplier of FHIR-native CDR software. Their peer-reviewed status and their grade are unchanged; what is narrower than it appears is their independence.

Case 1: DermaDashboard (oncological cohort building)

Borys K, Hartmann EM, Idrissi-Yaghir A, Livingstone E, Lodde G, Schmidt CS, Winnekens P, Friedrich CM, Hosch R, Nensa F (2025): *DermaDashboard: Bridging the Gap Between FHIR Standards and Clinical Usability.* JMIR Cancer 11:e73691. DOI: 10.2196/73691. *Evidence grade C (peer-reviewed viewpoint).*

Architecture: a relational FHIR store (commercial model, PostgreSQL-based), SQL queries without a REST detour, visualisation via an open-source tool. Use case: structured cohort filtering and exploratory data analysis for non-technical clinical users across several source systems, demonstrated on melanoma. The authors describe the architecture as transferable to other disease domains.

Limits: proof-of-concept character; no systematic performance evaluation; no comparison against an openEHR-based alternative; query performance on complex many-to-many queries not quantified; the FHIR engine is a commercial product, and transferability to other FHIR stores is not addressed.

Case 2: Real-time dashboard for blood-product management

Schmidt CS, Wutzkowsky J, Schäfer H, Reinartz L, Chahin H, Winnekens P, Alsulaiman JA, Temme C, Lenz V, Hosch R, Nensa F, Friedrich CM, Böckmann B, Horn PA (2026): *A real-time dashboard for optimizing blood product utilization through a FHIR-based data integration pipeline.* Transfusion and Apheresis Science 65:104301. DOI: 10.1016/j.transci.2025.104301 (online since 19 Dec 2025). *Evidence grade C (peer-reviewed short report).*

Architecture: an ETL pipeline from HIS and LIS into a FHIR store; a dashboard in daily clinical routine operation. Use case: real-time identification of blood products that can be released after a patient's discharge, transfer or death.

Significance: This case goes beyond analytical secondary use: the dashboard is integrated into an operational clinical workflow and used daily. It thereby supports the thesis that FHIR-native data holding is not confined to exchange scenarios.

Limits: no formal evaluation design; the use case is operational-logistical, not diagnostic, in nature; the ETL step between HIS/LIS and the FHIR store is a transformation layer. This is **not** native FHIR primary documentation.

D-7.4 Infrastructure Context

The German Medical Informatics Initiative has established FHIR as the implementation basis for the core dataset of all university data integration centres (Ammon et al., *Bundesgesundheitsblatt* 2024, 67:656–667, DOI: 10.1007/s00103-024-03888-4; evidence grade B).

At national aggregate level, the initiative's public research data portal (Forschungsdatenportal für Gesundheit, FDPG) displays the pooled holdings of the reporting data integration centres. As of 22 July 2026, across 41 centres, that is more than three billion laboratory values and over 385 million medication records. It is also some 350 million diagnoses and 150 million procedures, on structured data from more than 27 million persons.

Evidence grade C — a self-reported public aggregate of the coordinating body, continuously updated and publicly auditable in the portal, but not independently verified; the figures are approximate sums that vary with the reporting sites at query time.

The DIZ of University Hospital Essen operates a FHIR store in the billion-resource range, with data from around one million patients. The figure comes from statements of the Institute for Artificial Intelligence in Medicine (IKIM) itself. *Evidence grade D — self-reported, not independently verified; to be treated as an order of magnitude, not as a robust figure.*

D-7.5 Further Production Cases at Regional and National Scale

The cases in D-7.2 to D-7.4 come from the European regulatory sphere and from Germany. Four further deployments are recorded here because they raise the scale at which FHIR-native persistence is documented. Three of them lie outside the region this assessment otherwise examines (see Document 3, Purpose and scope). All four rest on presentations or reports by the implementing parties. None carries an independent evaluation, and none contains a comparison against an openEHR-based alternative architecture. Evidence grade D throughout.

- **Turkey — national repository (Common Data Model).** A presentation at DevDays 2025 (Namlı) reports a FHIR-based national repository holding some 38 billion FHIR resources on around 32 million patients, together with access times measured by the implementing party itself. Of the cases known here, this is the largest documented FHIR resource holding.
- **Lombardy, Italy — regional clinical data repository.** A presentation at DevDays 2025 (Manzelli) describes a FHIR-based regional clinical data repository covering around 10 million patients. Alongside Wales at national level, this is the largest documented regional case in Europe.
- **Singapore — Healthier SG.** A presentation at DevDays 2024 (Chai) describes the FHIR implementation for the national Healthier SG programme. The published material characterises the FHIR use, but does not permit a determination of whether persistence is FHIR-native in the sense of D-7.2. The case is therefore carried as a FHIR implementation at national programme scale, not as evidence for Option B.
- **Ukraine — national electronic health record.** Reported as FHIR-native storage with approximately 39 million patients and 420,000 users (Smithies 2026, personal communication), with the qualification, stated by the reporting source itself, that the REST interface is not a FHIR API. On the delimitation this appendix uses, that makes it a case of FHIR persistence but not a FHIR-native CDR in the Welsh sense, where persistence model and exchange model coincide. It is recorded for the persistence scale and is not counted towards Option B.

What these cases change, and what they do not. They answer one objection precisely: that FHIR-native persistence is documented only at institutional scale or in small national systems. Turkey and Ukraine put the resource and population counts an order of magnitude above the European cases in this appendix. Lombardy adds a regional case of a size comparable to a mid-sized European state.

What they do not change is the evidence position. All are self-presentations by implementing parties, and the symmetry note at the head of this appendix applies to them unchanged. What is shown is feasibility and operation at scale, not superiority. Two of the four do not support the claim about Option B and are expressly set aside for it, which is recorded here rather than left to the reader to infer. All four run in a FHIR-favourable direction and are graded accordingly.

D-7.6 Analytical Appraisal

1. **Primary documentation versus secondary use.** The Essen cases and Wales operate on ETL-transformed FHIR data, not on care documented natively in FHIR. Appendix D-7.7 records a publicly funded German programme aiming at native FHIR capture, and Appendix D-7.8 records a product for which the vendor declares native FHIR capture, unverified by any inspected artefact. Neither is a demonstrated comprehensive instance. The cases therefore establish FHIR-native persistence as a target model, not native FHIR primary documentation.
2. **Query depth.** Operational dashboards with a limited spectrum of resource types are a different use case from complex longitudinal cohort studies with a heterogeneous data profile.
3. **Representativeness.** Measured by FHIR infrastructure and informatics capacity, UH Essen is not a representative European reference hospital; Wales is a centralised state system without the federal fragmentation of Germany or Switzerland.

Consequence for the argument. The cases indicate a changing practice landscape; they do not establish structural superiority. They support the feasibility of Option B. The evidence remains asymmetric: more FHIR persistence cases were identified, while the reviewed openEHR evidence includes a longitudinal implementation study of a kind not found for the FHIR cases assembled here.

D-7.7 The German Vendor Market — Tiplu, and the Distance Between a Product and a Programme

This case qualifies the appraisal in D-7.6 but does not overturn it.

Tiplu GmbH, Hamburg, offers **TipluDB**, described by the vendor as a FHIR clinical data repository. Two propositions have to be kept apart here, because the material runs them together and they belong in different rows of § 1.6.1.

The product is Option B. The product page states that TipluDB "brings the patient and health data present in your institution **from the HIS and subsystems** into a uniform structure", and that the data of currently admitted patients are updated automatically and comprehensively. Master data, movement data, written documentation, billing data, laboratory data, vital parameters and medication are named as the scope.

That is capture in a proprietary primary system followed by persistence in a FHIR store: one crossing, persisted. **Evidence grade D** — vendor product documentation without independent evaluation, retrieved 5 August 2026.

What it adds that the earlier cases did not. Appendix D-7.2 to D-7.5 evidence Option B at national and infrastructural level — a health system, a research infrastructure, a university hospital. None of them evidences it as something a German institution can **buy**.

A commercially offered FHIR CDR in the German market is a different kind of evidence for the same architecture, and it bears directly on the procurement recommendation of § 1.6.3, which is addressed to buyers rather than to builders. It does not bear on effectiveness, and the symmetry note at the head of this appendix applies unchanged.

The programme aims at Option C, and has not reached it. The Federal Ministry of Health funded a research project, *Entwicklung eines intersektoralen FHIR-Datenmodells*, carried out by Tiplu and running to the end of February 2026, with co-financing under EU-HIP.

Its stated aim is that treatment data be "collected sector-independently **as FHIR data** and made available interoperably", enabling direct integration into the electronic patient record "without additional transformation". The company's managing director describes the approach as the consistent development of the medical information objects and "the structured collection of treatment data as FHIR data".

That is a description of Option C, and it is the first instance in the material examined for this assessment of Option C being pursued as a funded, named objective rather than posited as a structural floor. **Evidence grade C** for the fact and the framing of the funding, consistently reported across two outlets independent of the company; **evidence grade D** for every architectural claim, all of which trace to the company.

What it does not show, stated precisely. The prototype work is reported as running on anonymised records supplied by a university hospital through TipluDB, that is, on data extracted from primary systems, which is Option B. No result report of the project was available at the time of writing, and the project period has ended.

The proposition supported is therefore: *Option C is a publicly funded development objective in Germany, with a data model under development and prototypes built on extracted data.* The proposition not supported is: *a primary system exists that keeps its clinical record in FHIR.* The distance between those two sentences is the whole content of this subsection.

D-7.8 nursIT Institute — Option C on the Vendor's Own Account, Unverified

This case is carried as a candidate and is not counted as a demonstrated Option C implementation.

nursIT Institute offers **careIT**, described as a clinical workplace system for nursing, medical and therapy documentation **at the point of care**. That is a primary capture system rather than a downstream repository. The portfolio lists **Smile CDR** as the enterprise clinical data repository, described as "FHIR-native, ISiK-conformant, AI-ready", and announces **careIT Voice** as speech-based FHIR documentation.

careIT is stated to connect to existing hospital information systems via FHIR, "without migration, without a break in the system". Named customers include hospitals in Germany and a provider in Lower Austria.

Why it matters for the argument. If careIT writes its clinical content directly as FHIR resources into that repository, the arrangement is Option C, native FHIR capture, in the nursing domain, in the German-speaking market, with named installations. That is the row § 1.6.1 carries as having no independently verified instance.

The published material did not settle it. "Connects to existing hospital information systems via FHIR" is a statement about integration, not about persistence, and a system can hold a proprietary clinical model and synchronise into a FHIR repository without any of its clinical content being kept in FHIR. The sibling product SENSDOC points the other way: it is described as running on a **SENSDOC database** into which all nursing measures, observations and sensor data flow, with FHIR named for interoperability *between the system components*. On the distinction of § 1.6.1 that reads as Option A or B.

The question was therefore put to the vendor, and answered. The author wrote to the owner and the chief technology officer of nursIT Institute; the reply of 17 August 2026 states that the products are "100% native FHIR", that the web application and the native iOS and Android applications speak to the FHIR repository directly through the FHIR client, including by SMART on FHIR, and that there is "nichts dazwischen", nothing in between. Smile CDR is supplied where the customer has no repository; where one exists it is used instead, provided it supports FHIR R4, and eight hospitals are stated to run on a third-party repository on that basis. Repository independence is described as a deliberate strategy.

What that establishes, exactly. On the vendor's own account careIT is an instance of Option C: clinical content written directly as FHIR resources at the point of care, with no intermediate clinical model. **Evidence grade D** — a vendor statement about its own product, obtained by direct enquiry, with no independent evaluation and no artefact inspected. That is the same grade the published product pages carry (D, vendor web pages, retrieved 5 August 2026), and a reply to an enquiry does not raise it.

What the reply does not answer, recorded so that the gap stays visible. It does not address SENSDOC, whose published description names a database of its own, so the contradiction inside the portfolio is unresolved rather than removed. It names no product version, and no installation beyond the count of eight on one third-party repository. And it says nothing about how application state, workflow and derived indices are held, which is the part § 1.6.3 expressly places outside the requirement but expects a bid to describe.

The consequence for the taxonomy, stated narrowly. Row C moves from *no instance identified* to *one vendor declares it, unverified*. It does not move to *demonstrated*. What would move it there is an artefact rather than an account: an architecture description, a conformance record, or a published deployment description that a reader can inspect. The vendor adds that he is not aware of others taking this path and would be interested to know of any, which is an observation about his own market awareness and is carried as nothing more.

D-7.9 Phoenix — A Second Commercial Offering, and a Product Named "FHIR-Native" That Describes Itself as Option B

van Holt J, Nensa F, Sander A, Engels D, Müller A, Kaatze T (2026): *Datensouveränität im Krankenhaus: FHIR-basierte Clinical Data Repositories in der Praxis. Praxisleitfaden für Gesundheitseinrichtungen.* Whitepaper, DMI · ID Berlin · Firemetrics, 27 pages. *Evidence grade D: vendor documentation for a named product, no independent evaluation, authored by the supplying parties together with customer voices.*

Why a marketing document is admitted here at all. It is not admitted for anything it claims about benefit. It is admitted for what it says about its own architecture, because that runs against the product's own name and is therefore the one class of statement a vendor document can carry: an admission against the reading its title invites.

The product is called "Phoenix — das FHIR native CDR". Its own architecture description places it in Option B. The ingest layer "übernimmt Daten aus heterogenen Quellsystemen: KIS, LIS, PDMS, Archivsysteme, PDFs" across HL7 v2, FHIR and document interfaces, and "transformiert sie eingehende Daten vor der Speicherung in eine standardisierte FHIR-basierte Darstellung" (§ 1.2).

Clinical content therefore originates in primary systems, crosses one governed model boundary at ingest, and is then persisted in FHIR. That is the arrangement this document calls **Option B**, and the case studies confirm it: in both, data are "aus Primärsystemen" aggregated into the repository (§ 2.3.1, § 2.3.2).

The same pattern as tiplu, and the second instance is what makes it a pattern. D-7.7 recorded that a product marketed as FHIR-native was Option B on inspection. One instance is an anecdote about one vendor's word choice.

Two independent instances come from unrelated suppliers in the same market within the same year. They indicate that **"FHIR-native" is used in the German CDR market to mean that FHIR is the store, not that FHIR is where the clinical record is written**. That is the distinction § 1.6.1 draws between rows B and C, and the market usage collapses it. A procurement question that asks whether a product is "FHIR-native" will therefore not separate the two rows. The system-of-record test at § 1.6.3 will.

What the document adds to the German market picture. A second integrated commercial CDR offering, assembled by three named suppliers rather than by a hospital, with a stated operating model (CDRaaS qT). Two German deployments are named: an unnamed large university hospital (§ 2.3.1), and the Evangelisches Klinikum Bethel, part of University Hospital OWL. The Bethel repository is called ARCHE and is described as extending to affiliated hospitals and the ambulatory sector (§ 2.3.2).

Evidence grade D for all of it: none is independently reported, and the deployment descriptions are the suppliers' own.

What is expressly not taken from it. The opening figures (30% of the world's data generated in healthcare, 80% unstructured, 3% used for AI) carry no source and are not used anywhere in this document. The efficiency claims (66% of discharge-letter writing time saved; a CAUTI prediction model above 0.90 F1) are vendor-reported without method or denominator. The comparison table of interoperability platform against CDR (§ 1.3) is a positioning device, not a finding.

One statement is a camp position and is recorded as such. On openEHR the document states that it "bietet jedoch nicht denselben Grad an standardisierter API-Exposition, breiter internationaler Implementierungsreife oder die regulatorische Ausrichtung, die FHIR bietet" (§ 3.3).

The proposition may or may not be defensible, but the source cannot carry it: it is a supplier of the competing architecture writing about the architecture it does not sell. It is recorded here as an industry position and is used nowhere as evidence. The symmetry rule applies: this document treats a FHIR-side camp statement exactly as it treats the openEHR-side camp statements catalogued in Appendix E.

An interest that touches evidence already carried in this annex, disclosed because the rule requires it. The second author, Felix Nensa, is listed as "UK Essen | Firemetrics", that is, affiliated with both the deploying hospital and one of the three suppliers behind the product. He is also a co-author of Borys et al. (2025) and Schmidt et al. (2026), the two peer-reviewed Essen papers carried at D-7.3.

Those remain peer-reviewed and their grade is unchanged, but the independence implied by "peer-reviewed university hospital report" is narrower than it looked: an author of both papers holds a commercial interest in the class of architecture they describe. **This qualification runs against the direction this document argues**, since D-7.3 is FHIR-side evidence, and it is recorded for that reason rather than in spite of it.

D-8: Systematic Reviews of FHIR Use in Health Research

This appendix evaluates four secondary studies of differing methodology. They are a review per the Kitchenham methodology (Ayaz et al. 2021), a PRISMA review prospectively registered in PROSPERO (Vorisek et al. 2022), and two scoping reviews per the PRISMA-ScR extension (Tabari et al. 2024; Duda et al. 2022, which counts projects rather than publications). Only the second follows PRISMA in the strict sense, and the difference bears on what each of them can carry (D-8.6).

They have been included because they make testable a claim frequently advanced in the debate: that FHIR is an exchange format and structurally too shallow for research-grade clinical modelling.

D-8.1 Preliminary Remark — What These Sources Deliver and What They Do Not

They do not deliver: a comparison with openEHR. None of the four reviews includes openEHR as a comparator; openEHR is practically absent from them. They therefore **cannot ground any statement on the relative modelling depth** of the two standards. Whoever uses them that way repeats the error § 1.4.1 criticises in the Catalonia example: reading feasibility as superiority.

They likewise do not deliver: proof of effectiveness. They are literature reviews, not evaluation studies. They map what has been published, not what works.

They do deliver: a systematic, methodologically traceable inventory of the **breadth and depth of application** of FHIR in health research; no equivalent inventory of openEHR alone was found.

The one secondary study that takes the standards together is the 2026 cross-standard review assessed separately at Appendix E-9, and it counts eight openEHR studies against thirty-nine for FHIR. The asymmetry is therefore in the

volume of the literature rather than in the availability of a review, and it speaks in itself neither for nor against openEHR (Finding 3).

D-8.2 Tabari et al. (2024) — FHIR-Based Data Models

Tabari P, Costagliola G, De Rosa M, Boeker M (2024): *State-of-the-Art Fast Healthcare Interoperability Resources (FHIR)-Based Data Model and Structure Implementations: Systematic Scoping Review.* JMIR Med Inform 12:e58445. DOI: 10.2196/58445.

Methodology: Scoping review per PRISMA-ScR; searches in PubMed, Scopus, Web of Science, IEEE Xplore, ACM Digital Library, Google Scholar; no temporal restriction; 466 papers identified, 31 included.

Core findings:

- Two model types: **dynamic, pipeline-based** models (25 of 31, i.e. 81%) and static models (6 of 31). The dominance of dynamic models refutes the notion that FHIR is used primarily as a static exchange format.
- Application domains: chronic diseases, COVID-19 and infectious disease, cancer research, intensive care, unstructured clinical notes.
- Most frequent resources: `Observation` (68%), `Condition` (61%), `Patient` (58%), `Procedure` (52%).
- Limits named: data integration, interoperability, standardisation, performance, scalability/generalisability.

A finding that speaks against FHIR, and therefore stated here explicitly: the paper by Bennett et al. reported in the review (MIMIC-IV to FHIR) records a **considerable increase in storage requirements** on FHIR conversion. The HAPI FHIR format needed markedly more storage space than the relational source structure. This finding is incorporated in § 2.4 ("Limits that must be named in fairness").

Critique: Scoping reviews map; they do not appraise. No systematic quality assessment of the included papers is performed. The conclusion that FHIR is a "highly promising interoperability standard" is a summary of the included papers' self-statements, not an independent judgement. **Evidence grade B/C.** No interests declared.

D-8.3 Vorisek et al. (2022) — FHIR in Health Research

Vorisek CN, Lehne M, Klopfenstein SAI, Mayer PJ, Bartschke A, Haese T, Thun S (2022): *Fast Healthcare Interoperability Resources (FHIR) for Interoperability in Health Research: Systematic Review.* JMIR Med Inform 10(7):e35724. DOI: 10.2196/35724. PROSPERO: CRD42021235393.

Methodology: Systematic review per PRISMA, prospectively registered in PROSPERO; searches in PubMed/MEDLINE, Embase, Web of Science, IEEE Xplore, Cochrane Library; 998 papers identified, 49 included. **This is the methodologically strongest of the four reviews** — the pre-registration is a quality feature the other two lack.

Core findings:

PURPOSE OF FHIR USE	SHARE OF STUDIES
Data standardisation	41%
Primary data capture	29%
Cohort recruitment	14%
Analysis	12%
Consent management	4%

- Specialist domains: generic approach (55%), infectious disease (16%), oncology (12%), genomics (8%), plus environmental medicine, genomic cancer medicine, neuroimaging, pulmonary hypertension.

- **Transferability:** 12 of 22 domain-specific studies report their solution as transferable to other use cases.
- **Terminology binding:** SNOMED CT in 29%, LOINC in 37%, ICD-10 in 18%, OMOP CDM in 12% of studies. Relevant to § 2.5.1: external terminology binding is documented practice in the FHIR research environment.
- Germany contributes 47% of the studies — a direct effect of the Medical Informatics Initiative.

Two finds of independent value for this assessment:

1. **Rinaldi E, Thun S (2021), two publications of the same work.** The conference paper *From openEHR to FHIR and OMOP data model for microbiology findings*, Stud Health Technol Inform 281:402–406, and the longer journal article *Mapping from openEHR to FHIR and OMOP CDM to support interoperability for infection control*, GMS Med Inform Biom Epidemiol 17(2):Doc07, which carries an attachment with the element-by-element map. A documented openEHR → FHIR → OMOP transformation path; used in § 1.6 and § 2.4.1. **Only the journal article carries a quality assessment**, scored by the authors against ISO/TR 12300: equivalence 0 for openEHR to FHIR, no extension required and full element coverage for the dataset examined, against 0.9 for FHIR to OMOP with a good match for about 78% of elements. DiagnosticReport and ServiceRequest have no counterpart in OMOP and commentary fields go into a separate Note table, which the authors read as the expected consequence of OMOP's population-level purpose. It is used in § 2.7 for the one measurement that exists on the FHIR-to-OMOP side, and it is used with its weight stated: a self-assessment by the mapping team, called subjective by the authors, who score their own method of validation at 2 and their own documentation of decisions at 2, on a single microbiology dataset against FHIR R4 and OMOP CDM v6. **The 100% coverage from openEHR to FHIR is a second, separate use**: it is a concrete counter-instance to the blanket claim that FHIR cannot carry the depth of an archetype (§ 2.6), and it comes from an openEHR-based consortium.
2. **HiGHmed Data Sharing Framework** (Hund et al. 2021; Wettstein et al. 2021, two contributions): the framework coordinates distributed processes and exchanges information through FHIR R4 resources. These publications establish the FHIR-based process and sharing layer. They do not by themselves establish that every clinical payload follows an openEHR→FHIR path. A documented site-level mapping from openEHR to FHIR is needed before a particular implementation is classified as Option D (§ 1.6.1).

Critique, and a remarkable self-admission. The authors record that a quality assessment of the included papers was not possible for lack of established instruments for technical standardisation contributions. And further: they assume that the published literature depicts a surplus of successful FHIR initiatives, because unsuccessful ventures typically remain unpublished; their review may therefore suffer from **publication bias**. This openness increases the credibility of the work, and at the same time rules out its use as evidence of effectiveness.

Interests. The last author (S. Thun) is deputy chair of HL7 Germany; declared in the article. By the same standard that Appendix E-1 applies to the Tsafnat editorial, this must be disclosed. The difference from E-1 is considerable: here we have a pre-registered, peer-reviewed systematic review, there a non-peer-reviewed editorial. The interest situation reduces the weight; it does not devalue the methodology. **Evidence grade B.**

D-8.4 Ayaz et al. (2021) — FHIR Implementations and Challenges

Ayaz M, Pasha MF, Alzahrani MY, Budiarto R, Stiawan D (2021): *The Fast Health Interoperability Resources (FHIR) Standard: Systematic Literature Review of Implementations, Applications, Challenges and Opportunities.* JMIR Med Inform 9(7):e21929. DOI: 10.2196/21929. (Corrected version; correction notice: 10.2196/32869.)

Methodology: Systematic literature review per the Kitchenham methodology; 8,181 papers identified, 80 included; period 2012–2019.

Core findings: A taxonomy of FHIR implementation models, resources used (around 82 different ones), favoured application types, and migration approaches from HL7 v2/CDA to FHIR. Challenges named: implementation complexity, standard complexity, adoption, maintenance, mapping/migration.

A finding relevant to this assessment: the authors conclude that FHIR will supersede the existing HL7 standards (v2, v3, CDA) **not within months but over years to decades** — worldwide, HL7 v2 is too deeply entrenched. That supports the market-inertia thesis (Finding 5) and the description of architecture Option A (§ 1.6.1), independently of the openEHR question.

Critique: the weakest of the four sources. The quality assessment is purely structural (checking, for instance, whether a paper states its research question and draws a conclusion); **not a single paper was excluded for lack of quality**. The closing recommendation that FHIR is "a future suitable solution" for the interoperability problem is a judgement that a review without a comparator standard cannot carry. Evidence grade C. Usable as an inventory, not as an argument.

D-8.5 Duda et al. (2022) — FHIR in Clinical Research, Counted by Project Rather Than by Paper

Duda SN, Kennedy N, Conway D, Cheng AC, Nguyen V, Zayas-Cabán T, Harris PA (2022): *HL7 FHIR-based tools and initiatives to support clinical research: a scoping review.* J Am Med Inform Assoc 29(9):1642–1653. DOI: 10.1093/jamia/ocac105.

A PRISMA-ScR scoping review that differs from the three above in what it counts. Instead of publications it counts **projects** — papers, funded initiatives and tools together — screening 1,572 candidates down to **203 FHIR projects applicable to clinical research**.

The distribution is the reason it is carried here: data mapping to and from FHIR 66 (32.5%), ontology management 30 (14.8%), sharing through repositories or registries 24 (11.8%), general research data sharing 23 (11.3%), genomic data 21 (10.3%), phenotyping 19 (9.4%). The authors' own reading is that activity concentrates on preparing data and connecting systems, and thins out inside the research process itself.

What it adds and what it does not. It corroborates D-8.2 to D-8.4 on breadth of use and adds nothing to the modelling-depth question, having no comparator standard either. Its distinct contribution is the one figure that a publication count cannot produce: **the single largest category of FHIR work in clinical research is mapping data into and out of FHIR**. That is a statement about where the effort goes rather than about what the standard can do.

Two authors hold roles at HL7 Da Vinci and the National Library of Medicine, and the work was funded under an NLM contract to promote FHIR for research. That is declared in the paper, and it is a plain interest in the subject. **Evidence grade B/C**: a systematic search with a documented protocol, counting self-described project purposes that were not independently verified.

D-8.6 Summary Appraisal

	AYAZ 2021	VORISEK 2022	TABARI 2024	DUDA 2022
Methodology	Kitchenham SLR	PRISMA + PROSPERO	PRISMA-ScR	PRISMA-ScR
Included	80	49	31	203 projects
Pre-registered	no	**yes**	no	no
Quality assessment	purely structural	not possible (disclosed)	not systematic	not applicable
openEHR as comparator	no	no	no	no
Evidence grade	C	**B**	B/C	B/C

What can be said on this basis: the claim that FHIR is structurally too shallow for research-grade clinical modelling does not hold as a blanket statement. FHIR is used productively in oncology, genomics, infectious disease and imaging for capture, standardisation and analysis, not only for exchange.

Evidence grade B for statements on breadth of application, carried chiefly by Vorisek; C to B/C for the others; **not usable** for effectiveness or superiority claims.

What cannot be said on this basis: that FHIR is superior to openEHR in modelling depth, or that FHIR-based research infrastructures are more effective. Both would require a comparator standard and an evaluation design, which none of the four provides. The evidence gap on *effectiveness* remains for both standards (Finding 3, Finding 4).

D-9: Bibliometric Comparison — Publication Volume over Time

In four lines, for a reader who does not need the method. *Question:* how much scholarly publication names each standard, and how has that changed. *Method:* title-and-abstract queries in OpenAlex and PubMed, a baseline and six sensitivity checks, twenty-one queries in all, recorded and reproducible.

Result: FHIR crossed over around 2016 and grew steeply while openEHR flattened and then declined. The annual gap is roughly ten to one against totals of roughly four to one, so the divergence is recent and widening. What widens it in relative terms is openEHR's decline rather than an acceleration specific to FHIR.

Limits: this counts research attention and nothing else — not deployment, not interoperability, not outcome — and one of the sensitivity checks shows the ratio varies by a factor of more than forty across document types. **A reader who accepts those four lines can skip to D-10.** The remainder sets out the method and every check, including the two that were proposed and not carried out.

This appendix quantifies the "adoption, implementation density and ecosystem maturity" lead noted in Finding 3. It measures one thing only, the volume of scholarly publication that names each standard, and it is included because that volume is a reproducible, if indirect, indicator of research attention and ecosystem momentum. The same distinction the document draws throughout (regulatory success is not empirical benefit, § 1.4/Finding 4; Appendix D-4 preliminary remark) applies here in full.

Method. Two open bibliographic databases were queried, restricting the search term to title and abstract. In OpenAlex the filter `title_and_abstract.search` was used with the terms `FHIR` and `openEHR`, grouped by `publication_year`. In PubMed the queries `FHIR[tiab]` and `openEHR[tiab]` were run as a conservative cross-check. OpenAlex counts more generously than PubMed because it indexes conference papers, preprints, software and datasets in addition to journals.

The first retrieval was made on 23 July 2026. A second retrieval on 2 August 2026 repeated the baseline and added six sensitivity checks, twenty-one queries in all. The figures below are those of 2 August 2026 unless a date is given, and both dates are shown wherever they differ. The six checks are set out in D-9.1 to D-9.5, D-9.3 carrying two of them, and each is reported whether or not it supports the finding.

Totals.

DATABASE	FIELD	FHIR	OPENEHR	RATIO
OpenAlex	Title/abstract, 2 Aug 2026	3,505	837	4.2 : 1
OpenAlex	Title/abstract, 23 Jul 2026	3,421	837	4.1 : 1
PubMed	Title/abstract `[tiab]`, 2 Aug 2026	1,061	291	3.6 : 1

The ratio is stable across two independently maintained databases with different coverage, and across a ten-day interval. How far that stability extends to other ways of measuring the same thing is the subject of D-9.1 to D-9.5.

How far the counts move between retrievals. Over ten days the OpenAlex FHIR total rose by 84 records, 79 of them in the partial year 2026. The openEHR total did not move at all and stood at exactly 837 on both dates. Beneath that unchanged total the annual distribution did move: 2015 and 2019 each gained one record, 2026 gained one, and 2023 lost three. The FHIR series likewise lost three records in 2015 and one each in 2018 and 2019 while gaining elsewhere.

Retrospective years are therefore not fixed, and the counts are not reliable to the unit in either direction. What the appendix rests on is the order of magnitude and the shape.

Annual counts (OpenAlex, title/abstract, 2 August 2026).

YEAR	FHIR	OPENEHR
2013	3	48
2014	12	47
2015	40	58
2016	75	60
2017	95	60
2018	137	41
2019	180	41
2020	176	45
2021	216	35
2022	238	38
2023	269	31
2024	353	36
2025	689*	52*
2026	985*	29*

*2025 and 2026 are provisional. 2026 is a partial year (query date 2 August 2026) and already exceeds the whole of 2025 for FHIR, with five months of the year still to run. That is the clearest single indication of how far OpenAlex over-counts very recent years: online-first records and preprints are indexed early. The upward direction is robust; the exact magnitude of the last two points is not. PubMed reports these years more conservatively. Records dated before 2013 are not shown: 216 for openEHR and 37 for FHIR, the latter mis-dated noise.

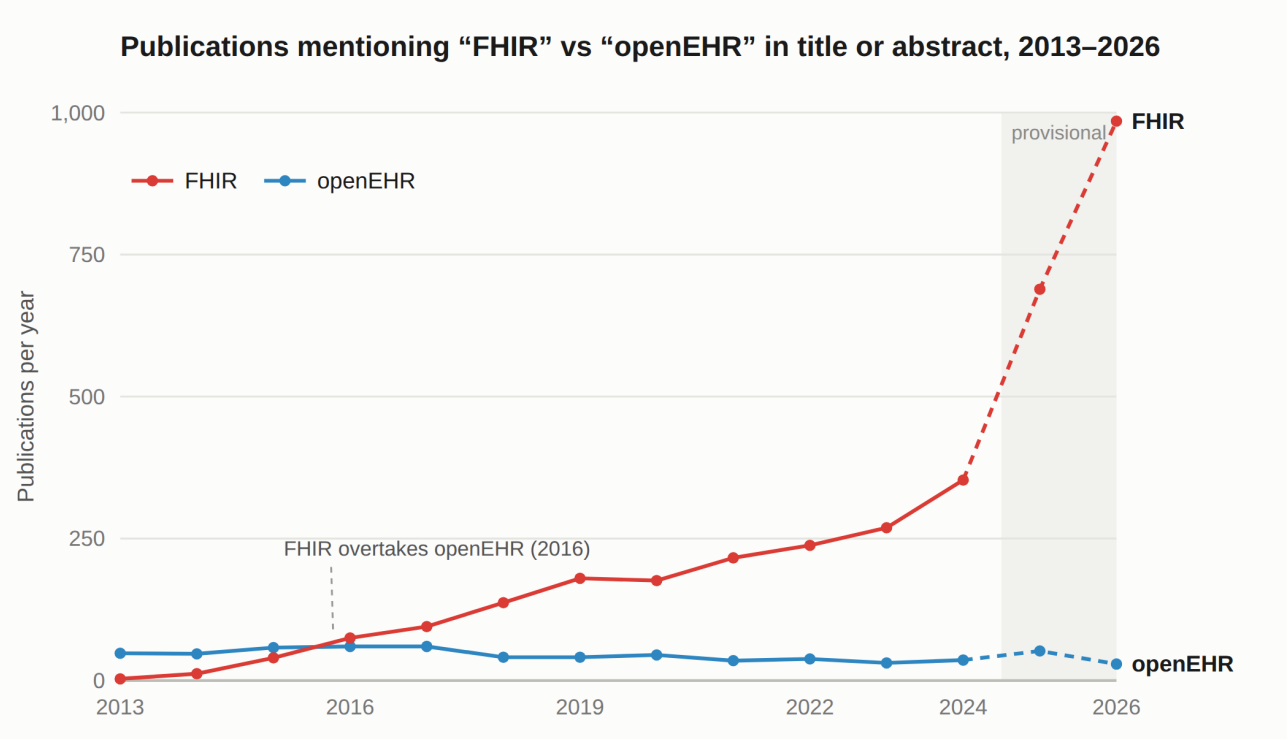

Figure 5 — Publications mentioning "FHIR" vs "openEHR" in title or abstract, 2013–2026. Source: OpenAlex, query 2 August 2026 (title/abstract search). Total records: FHIR 3,505; openEHR 837. Cross-check PubMed [tiab]: FHIR 1,061; openEHR 291. Shaded 2025–2026 provisional (OpenAlex over-counts recent years; 2026 partial). Counts reflect research attention, not interoperability or effectiveness.

D-9.1 How far the finding depends on the search string

Three retrievals per standard, all on 2 August 2026.

SEARCH	FHIR	OPENEHR	RATIO
Title/abstract, acronym (the published query)	3,505	837	4.2 : 1
Full search, including indexed full text	11,698	2,418	4.8 : 1
Narrowest string (see below)	1,140	34	not comparable

The two comparable tiers give 4.2 : 1 and 4.8 : 1, and PubMed gives 3.6 : 1. Across these three measurement choices the ratio lies between roughly 3.6 : 1 and 4.8 : 1. Widening the search from title and abstract to indexed full text roughly triples both counts and leaves their relation intact. The finding therefore does not depend on the restriction to title and abstract.

The third row does not compare, and the reason is a fault in the design of the check rather than in the data. It was intended as a pair of narrowest plausible strings, but the two are not counterparts. For FHIR it is the expanded name `"Fast Healthcare Interoperability Resources"` as an exact phrase, for openEHR the two-word spelling `"open EHR"`. openEHR has no expanded name, so no equivalent of the first exists for it. The mismatch is recorded rather than repaired, because the two counts remain informative taken separately.

The expanded FHIR name appears in the title or abstract of 1,140 records, 32.5% of the acronym corpus. That puts a floor under the share of that corpus which is unambiguously about the standard. The two-word openEHR spelling returns 34 records. One work in the random sample of D-9.2 carries "Open EHR" in its title and is also present in the closed-spelling corpus. The two sets therefore overlap, and the 34 cannot be added to the 837.

D-9.2 What the records are: a title-level review of two random samples

Fifty records were drawn at random from each corpus with a fixed seed (`sample=50&seed=42`), so that the draw can be repeated by anyone with access to the database. Four fields were requested: identifier, title, publication year and type. The classification is therefore a review of titles, not of abstracts or full texts, and that is its principal limit. A title can establish that a record is out of scope; it cannot establish what a record whose title does not name the standard is about.

Three classes were fixed before the titles were read. **Out of scope**: the record is not about health care, health data or health information at all, or it is not a work. **Named**: the standard appears in the title, so it is a subject of the work or its declared instrument. **Not named**: the record is in scope and its title does not name the standard, which a title-level review cannot resolve further.

	FHIR	OPENEHR
Out of scope	2	0
In scope, standard named in the title	17	28
In scope, standard not named in the title	31	22

The two out-of-scope FHIR records are a paper on quality assurance in high-frequency trading systems and a record whose entire title is an HTML fragment reading "Data flow and technical infrastructure at ACCB", indexed as a work of type `other`. The second is an indexing artefact rather than a false hit on the search term.

A further four FHIR records and two openEHR records could not be placed from the title alone, among them a work on pharmaceutical manufacturing quality data, an essay on programmable built environments and a book chapter titled "Standards Development Organizations". Read as a band, the false-positive rate is 4% to 12% for FHIR and 0% to 4% for openEHR.

This supports the claim that both terms are specific acronyms carrying little false-positive risk, and supports it more strongly for openEHR than for FHIR. It also shows what a count of mentions does not show.

Only 17 of the 50 FHIR records name the standard in the title, against 28 of the 50 openEHR records. The two corpora are therefore not the same kind of corpus. A larger part of the FHIR literature names the standard somewhere in an abstract while being about something else, and a larger part of the openEHR literature is about openEHR.

At fifty records per sample this difference sits at the edge of what the sample can distinguish, and it fixes a direction rather than a size. It is nonetheless the reason the totals must be read as attention and not as a body of work on the standard.

Two further observations from the same samples. The oldest FHIR record drawn is from 2018 and the oldest openEHR record from 2005, with nine openEHR records before 2013, which is consistent with the first observation below. And one record appears in both samples, on cross-institution health data sharing in Norway using both standards, so the two corpora overlap. The size of the overlap was not measured.

D-9.3 Document types, and reviews separated from the rest

TYPE	FHIR	SHARE	OPENEHR	SHARE	RATIO
Journal article	1,471	42.0%	336	40.1%	4.4 : 1
Book chapter	507	14.5%	157	18.8%	3.2 : 1
Conference paper	451	12.9%	156	18.6%	2.9 : 1

TYPE	FHIR	SHARE	OPENEHR	SHARE	RATIO
Preprint	430	12.3%	42	5.0%	10.2 : 1
Software	181	5.2%	3	0.4%	60.3 : 1
Doctoral thesis	108	3.1%	83	9.9%	1.3 : 1
Dataset	88	2.5%	16	1.9%	5.5 : 1
Other	81	2.3%	25	3.0%	3.2 : 1
Conference abstract	53	1.5%	5	0.6%	10.6 : 1
Report	29	0.8%	2	0.2%	14.5 : 1
Review	28	0.8%	3	0.4%	9.3 : 1
Book	22	0.6%	3	0.4%	7.3 : 1

Shares are of the respective corpus. The listed types account for 3,449 of the 3,505 FHIR records and 831 of the 837 openEHR records; the remainder carry types below the reporting threshold or none.

The two corpora carry a similar share of journal articles, 42.0% against 40.1%. Restricting the comparison to journal articles leaves the ratio at 4.4 : 1, close to the headline figure. Otherwise the ratio by type varies by a factor of more than forty, from 60 : 1 for software to 1.3 : 1 for doctoral theses.

The spread has a readable direction. The FHIR lead is widest in the fastest-moving and least reviewed categories, software and preprints, and narrowest in the slowest, doctoral theses and conference papers. That is the profile of a lead in current momentum. It is not the profile of a lead in accumulated, reviewed evidence. The headline ratio of the totals should be read as an average across categories that disagree, rather than as a single stable quantity.

Reviews were separated from the rest by the PubMed publication type `review[pt]`: 94 of the 1,061 FHIR records, 8.9%, and 7 of the 291 openEHR records, 2.4%. OpenAlex's own `type` field classifies 28 FHIR records, 0.8%, and 3 openEHR records, 0.4%, as reviews. The two classification systems disagree by roughly an order of magnitude in the share.

That is a statement about the instruments rather than about the literature. The OpenAlex type field is not a usable review indicator at this granularity, and only the PubMed figure is relied on here. On that figure the FHIR literature carries proportionally about three and a half times as many review articles as the openEHR literature.

Whether the remainder are primary implementation studies is not established by either database, because neither carries a field separating a primary implementation study from a descriptive one. That separation would require reading the records, which was not done.

D-9.4 Where the work comes from

Country here means the country of an author's affiliated institution, not the country in which a system was deployed or studied. A work with authors in several countries counts in each of them, so these figures are not shares of a partition and do not sum to the corpus.

RANK	FHIR CORPUS	RECORDS	OPENEHR CORPUS	RECORDS
1	United States	868	Germany	88
2	Germany	444	Portugal	64
3	United Kingdom	139	Brazil	59

RANK	FHIR CORPUS	RECORDS	OPENEHR CORPUS	RECORDS
4	India	136	Spain	56
5	Austria	117	China	43
6	Italy	106	United States	40
7	Spain	102	United Kingdom	35
8	Canada	94	Sweden	31
9	Portugal	90	Netherlands	26
10	Australia	81	Australia	24

The structures differ. One country dominates the FHIR corpus: the United States accounts for 868 records, 24.8% of it. No country dominates the openEHR corpus, whose leaders are Germany with 88 records, 10.5%, then Portugal, Brazil and Spain; the United States is sixth with 40 records, 4.8%. Germany leads the openEHR corpus and is second in the FHIR corpus with 444 records, 12.7%.

That last figure bears on the third observation below. Vorisek et al. (Appendix D-8.3) report that Germany contributes 47% of FHIR research studies. The base is not the same. Their share is of the studies that passed the screening of a systematic review. The 12.7% is the share of all records whose title or abstract names FHIR that carry at least one German affiliation.

The two figures are not in contradiction, and the distance between them shows how strongly an apparent German concentration depends on which population is counted. The concentration claim is therefore retained on the base on which it was published and is not extended to the full corpus.

Restricting both corpora to European affiliations gives 1,152 FHIR records, 32.9% of that corpus, and 367 openEHR records, 43.8%. The European ratio is 3.1 : 1, narrower than the global 4.2 : 1: openEHR is proportionally more European than FHIR. The European series has the same shape as the global one, and the crossover falls in the same year.

YEAR	FHIR (EUROPEAN AFFILIATION)	OPENEHR (EUROPEAN AFFILIATION)
2013	0	27
2014	2	18
2015	6	30
2016	18	15
2017	23	20
2018	39	17
2019	54	15
2020	51	22
2021	75	20
2022	88	24
2023	107	10
2024	164	21

YEAR	FHIR (EUROPEAN AFFILIATION)	OPENEHR (EUROPEAN AFFILIATION)
2025	222*	28*
2026	300*	16*

*Provisional, as in the global series. The annual figures sum to less than the totals given above because works dated before 2013 are not shown: 84 for openEHR and 3 for FHIR.

D-9.5 Normalised against the growth of the field

No canonical corpus of digital-health publication exists, so no single denominator can be the right one. Three title-and-abstract phrases were therefore taken as denominators: "digital health", "health informatics" and "electronic health record". What is reported is the range across them rather than one normalised curve. None of the three contains either corpus as a subset, since a record naming FHIR need not carry any of the three phrases. The figures below are accordingly ratios between two independently defined counts, not proportions of a whole.

Between 2013 and 2024, the last year that is not provisional, "digital health" grew from 118 records a year to 5,815, a factor of 49. "Electronic health record" grew from 2,150 to 11,074, a factor of 5.2, and "health informatics" from 454 to 721, a factor of 1.6.

One of the three fails outright, and the failure is reported rather than dropped. In 2026 the FHIR count exceeds the "health informatics" count, 985 against 680, which a denominator cannot do if it is to stand for the surrounding field. The phrase itself fell out of use over the period. It measures a vocabulary, not a literature, and it is not used further.

YEAR	FHIR / "DIGITAL HEALTH"	OPENEHR / "DIGITAL HEALTH"	FHIR / "ELECTRONIC HEALTH RECORD"	OPENEHR / "ELECTRONIC HEALTH RECORD"
2013	2.5%	40.7%	0.1%	2.2%
2016	18.1%	14.5%	2.3%	1.8%
2019	12.3%	2.8%	3.2%	0.7%
2022	6.7%	1.1%	3.0%	0.5%
2024	6.1%	0.6%	3.2%	0.3%
2025*	6.3%	0.5%	4.6%	0.3%
2026*	8.9%	0.3%	8.2%	0.2%

Two things follow, and the second qualifies a word used elsewhere in this appendix.

Normalisation cannot alter the ratio between the two standards, because both are divided by the same denominator. The divergence between FHIR and openEHR is therefore untouched by this check, and it is real: openEHR's share falls against every denominator, continuously, from 2015 onward.

What normalisation does alter is the reading of FHIR's own curve. Measured against "digital health", the FHIR share rose to about 18% in 2015 and 2016 and has fallen since, standing at 6.1% in 2024 and roughly level from 2022. Measured against "electronic health record", it rose to about 3% by 2018 and remained there through 2024. On both usable denominators FHIR has grown at approximately the rate of the surrounding field since about 2019 rather than faster than it.

Much of the steepness of the absolute curve is the steepness of digital-health publishing as a whole. Both shares rise again in 2025 and 2026, but those are the years this appendix already marks as unreliable, and the rise is not read as a change of direction.

D-9.6 Reading the Finding, and Its Limits

Reading the finding. Three observations follow, each with its qualification.

First, openEHR is the older field. Before 2013 some 216 openEHR-naming works are recorded against effectively none for FHIR (the 37 pre-2013 OpenAlex hits for "FHIR" are mis-dated noise, which incidentally shows the raw counts are not reliable to the unit). In 2014 and 2015 openEHR still published more per year than FHIR.

Second, FHIR crossed over around 2016 and then grew steeply in absolute terms, while openEHR stayed flat and then declined. By 2024 the annual gap, roughly ten to one, is far wider than the ratio of the totals, about four to one: the divergence is recent and it continues to widen. Measured against the growth of digital-health publication as a whole (D-9.5), the FHIR share has been roughly constant since about 2019, so what widens the gap in relative terms is openEHR's decline rather than an acceleration specific to FHIR.

Third, the absolute acceleration tracks the regulatory timeline, not an effectiveness result. The FHIR literature is concentrated in mandated and publicly funded systems. Vorisek et al. (Appendix D-8.3) find Germany alone contributing 47% of FHIR research studies, on a screened study population rather than on the full corpus (D-9.4). That makes a programme effect of the Medical Informatics Initiative plausible without establishing one.

Two counts moving together do not settle which moved the other, and no attempt is made here to settle it. Read with that reservation, the bibliometric curve is consistent with Finding 4 (the decisive factor is regulatory) as much as with Finding 3. It is a measure of momentum and attention, and belongs to the same evidentiary class as adoption and tooling maturity, not to the class of effectiveness evidence, which remains thin for both standards.

Evidence grade B for the counts and ratios, which are recorded, reproducible queries (Appendix H-4); no grade for any proposition about deployment, interoperability or effectiveness.

Limits. Publication counts measure research attention, not deployment, interoperability or clinical outcome. Database coverage differs, so absolute counts are database-specific. The ratio and the temporal shape are the parts that hold across sources, and D-9.3 shows that even the ratio is an average over document types that disagree by a factor of more than forty.

The recent-year inflation in OpenAlex is real and flagged. The sample review of D-9.2 is a review of titles, not of abstracts, so it bounds the false-positive rate but does not establish how many of the remaining records treat the standard substantively. Country is the country of an author's institution and not of a deployment.

Two checks that were proposed and not carried out are named here rather than passed over. The corpora were not read for the distinction between primary implementation studies and descriptive ones, since neither database carries a field for it. And the overlap between the two corpora was not measured, although the samples show that it exists.

No systematic bibliometric inventory weighted by study quality was undertaken. This is a volume count, deliberately simple and reproducible, now with its sensitivity to the main measurement choices set out rather than assumed.

D-10: German FHIR Profile Governance — Published Separately

Appendix D-10 is published as **Document 5 — Technical Report: German FHIR Profile Governance**. It contains the corpus definition, package and dependency analysis, constraint counts, terminology-binding analysis and reproducibility information. It is not duplicated here.

D-11: Three Primary Studies — What Certification Delivers, What a Change of System Costs, and What "Same Vendor" Buys

This appendix holds three peer-reviewed primary studies, the first two added on 6 August 2026 and the third on 18 August 2026. They are grouped because each corrects a different possible over-reading. D-11.1 examines semantic errors and permitted variation in documents produced by certified technologies. D-11.2 examines changes in alert behaviour after one electronic record system replaced another. D-11.3 measures how far sites on the same or different vendor products recognise each other's representations of common structured elements. None of the three studies compares architectures or standards, and neither measures patient harm.

How they were found, disclosed because it bears on selection. Both were reached through a narrative review of transitions between electronic health records (Huang C, Koppel R, McGreevey JD III, Craven CK, Schreiber R, *Transitions from One Electronic Health Record to Another: Challenges, Pitfalls, and Recommendations*, Appl Clin Inform 2020;11:742–754, DOI 10.1055/s-0040-1718535). That review is **not drawn on for any proposition here** and carries no ledger row.

Three reasons were decisive and are recorded so that the decision can be challenged. Its own contribution is the authors' consensus rather than a finding. Its arresting cost figures are sourced to trade press rather than to a collection. And its unit of analysis is the replacement of a whole system, which is not the operating path this document is concerned with. It served as a finding aid and is named as one.

A caution that applies to both. Neither study compares architectures, and neither is about openEHR or FHIR. They are evidence about certification regimes and about system replacement respectively. Where they are cited in the body, the proposition they carry is stated narrowly for that reason.

D-11.1 D'Amore et al. (2014) — What Documents from 21 Certified Technologies Contained

D'Amore JD, Mandel JC, Kreda DA, Swain A, Koromia GA, Sundareswaran S, Alschuler L, Dolin RH, Mandl KD, Kohane IS, Ramoni RB (2014): *Are Meaningful Use Stage 2 certified EHRs ready for interoperability? Findings from the SMART C-CDA Collaborative.* J Am Med Inform Assoc 21(6):1060–1068. DOI: 10.1136/amiajnl-2014-002883. Open access. *Evidence grade B for the inspected C-CDA documents and reported observations; no grade for any proposition about FHIR, architectures or patient harm.*

Method. 107 health information technology organisations were invited, 44 responded, and 91 C-CDA documents from 21 distinct technologies were collected. All 91 were parsed automatically. **The manual inspection, which produced the headline count, covered 21 documents — one per vendor application.**

An *error* was defined as XML usage conflicting with mandatory guidance in the HL7 C-CDA 1.1 Implementation Guide; *heterogeneity* meant variation in inclusion, omission or expression that the specification permits. Two validators were run alongside the manual pass, NIST's Transport Testing Tool and the SMART C-CDA Scorecard.

Findings. 615 observations across those 21 documents, 607 of them assigned to one of six mutually exclusive categories:

CATEGORY	OBSERVATIONS
Element optionality through inclusion or omission	161
Terminology misuse or omission	142
Inappropriate or variable XML organisation or identifiers	110
Incorrect data within XML elements	97
Inconsistent data representation	52

CATEGORY	OBSERVATIONS
Problematic reference to narrative text from structured body	45
Not elsewhere classified	8
Total	**615**

Twelve distinct patterns were observed for recording a telephone number. On the SMART C-CDA Scorecard the 21 samples averaged 63%, ranging from 23 to 100. Of the 21 samples run through the NIST validator, ten returned no errors and the other eleven averaged 71, ranging from 2 to 297.

The finding that carries furthest is not a number. Several vendors explained to the investigators that they had focused development on generating documents that pass the validator, and less on the semantic interoperability implementers were asking of them. The authors' own recommendations follow from it. Reduce optionality, because the Common Data Set does not constrain it at a granular level and thereby permits omission. Require vendors to populate data they hold. And make conformance testing of the reference vocabularies part of certification.

Limits, and they are substantial.

1. **The 615 come from 21 documents, not from 91.** The distinction is easy to lose because both numbers appear in the abstract, and losing it overstates the sample while understating the density.
2. **Only the first instance of each problem type in each domain of each document was recorded.** The authors state the reason: counting repeats would measure the frequency of errors in one document rather than the prevalence of error types across the field. The number therefore understates frequency by design and should never be read as an error rate.
3. **The documents are fictional patient records exported from vendor test and development environments**, not production documents. The authors expect production exports to show the same problems and say they anticipate further ones, but that is their expectation and not their result.
4. **Participation was voluntary**, so the 21 technologies are self-selected among those willing to be examined.
5. Seven data domains of the Common Data Set were scrutinised; the authors state that examining others would have found more.
6. **C-CDA, and 2014.** The document standard, the certification regime and the market have all moved. What transfers is the mechanism, a permissive specification plus a validator that can be satisfied without semantic agreement, and not the figures.

Interests. Funded by an ONC SHARP award. Authors are drawn from Lantana Consulting, Diameter Health, Harvard and Boston Children's; several hold positions in organisations that consult on or supply document-quality tooling, and one author contributed to the open-source parser used. The direction of the finding runs against the certification regime the authors' field operates within rather than towards a commercial interest of theirs, which is recorded as an interest observation and not as a grade uplift.

D-11.2 Wright et al. (2018) — What Was Lost That Was Never in the Data

Wright A, Aaron S, Seger DL, Samal L, Schiff GD, Bates DW (2018): *Reduced Effectiveness of Interruptive Drug-Drug Interaction Alerts after Conversion to a Commercial Electronic Health Record.* J Gen Intern Med 33(11):1868–1876. DOI: 10.1007/s11606-018-4415-9. *Evidence grade A for what it measured at the site it measured; no grade for any proposition about architectures or standards.*

Method. A before-and-after study at Brigham and Women's Hospital, which replaced its internally developed record (LMR) with a commercial product on 31 May 2015 in a single cutover. Outpatient medication orders and alert events were extracted for the last six months of the old system and the first six of the new, covering 3,277 prescribers who

received at least one interaction alert. Acceptance was defined as cancelling the proposed order or discontinuing the interacting drug.

Findings.

MEASURE	OLD SYSTEM	NEW SYSTEM
Interruptive interaction alerts per 100 orders	about 2	about 11
Acceptance, interruptive alerts overall	30.7%	6.2%
Acceptance, most severe tier	100% (n = 304)	8.4% (n = 1,753)
Acceptance, medium tier	29.3% (n = 15,041)	7.5% (n = 38,074)

The 100% is an artefact and must be carried with the figure. In the old system the most severe combinations could not be ordered at all: a hard stop, not a warning. There was nothing to accept or decline, so the rate is a property of the configuration rather than a measure of clinical diligence. Quoting the fall from 100 to 8.4% as a decline in prescriber behaviour misreads it.

Two rival explanations were tested and largely excluded. The commercial product carries a larger interaction knowledge base, which would depress the average acceptance rate by adding less important alerts. But acceptance also fell for drug pairs present identically in both systems, among them pairs an expert panel had classed as never to be co-prescribed.

Alert fatigue was tested by a natural experiment. When the least severe tier was switched off on 1 March 2016, the interruptive burden fell by 50.5%. Acceptance of the most severe tier rose only from 9.1 to 12.7% and of the medium tier from 6.9 to 9.2%, both statistically significant and both small.

What the authors conclude. The findings are consistent with an important role for workflow and configuration: reduced differentiation between tiers, movement of the alert from drug selection to order signing and the ability to dismiss several alerts together. The observational design does not isolate any one feature as the sole cause. An earlier study at the same site supports the relevance of tiering, but does not turn the before-and-after comparison into a controlled architectural experiment.

Limits, as the authors state them. One site, one installation of one commercial product, which they expressly decline to generalise to other installations of the same product. Outpatient orders only. The first year after go-live, with no claim about a later steady state. Individual interface features not separated experimentally. And, decisively for how far it can be carried here, **no measurement of adverse drug events and therefore no statement about patient harm**. Some artefacts in the data are reported as unexplained.

Why it is carried, and how far. The study is not about data migration. It is the clearest measurement held here that a safety function can fail to survive a change of system even where the data arrive intact. The function lived in the configuration rather than in the record: tiering, hard stops, the point in the workflow at which the check fires, and whether an answer is required item by item.

No clinical information model on either side of the debate carries these properties. That is the whole of what it is cited for at § 1.6.2, and it cuts equally against over-claiming for openEHR persistence and for FHIR persistence.

Interests. Funded by the National Library of Medicine (R01LM011966) and AHRQ. No conflicts declared. The finding is unflattering to a commercial product the authors' institution had itself procured.

D-11.3 Bernstam et al. (2022) — A Number for What "Same Standard, Same Vendor" Delivers Between Sites

Bernstam EV, Warner JL, Krauss JC, Ambinder E, Rubinstein WS, Komatsoulis G, Miller RS, Chen JL (2022): *Quantitating and assessing interoperability between electronic health records.* J Am Med Inform Assoc 29(5):753–760. DOI: 10.1093/jamia/ocab289. Read in full. **Evidence grade B for the measured scores; no grade for any proposition about FHIR, openEHR or information loss, none of which the study examines.**

What was measured. Using de-identified data from 68 US oncology sites on 5 EHR vendor products in the ASCO CancerLinQ platform, the authors defined an interoperability score: the probability that an original *name string* for one of 12 data elements — six medications with RxNorm codes, six laboratory tests with LOINC codes — sent from one site had already been seen at the receiving site. The mean score was **0.68 within the same vendor's product and 0.22 between vendors** when weighted by the number of systems, 0.57 and 0.20 unweighted. Even hematocrit, a standard element with a well-accepted LOINC code, arrived under many distinct local representations.

What it supports here. It is the quantitative counterpart of § 1.5: sites that implement the same certified product of the same vendor are, on this measure, only about two-thirds mutually intelligible for common structured elements, and sites on different vendors about one-fifth, although recommended national standards (LOINC, RxNorm) exist for every element measured. The authors' closing sentence carries the finding of § 1.5 in their own words: "vendor choice does not ensure reliable interoperability. Reliable interoperability requires institutions to map their data to the same standards and ensure that mapping practices are consistent across institutions." That is a governance conclusion drawn from measurement, and it applies to any standard, which is why it is carried here without a side.

Boundaries, stated before anyone leans on the number. The measure is a string-recognition proxy: it asks whether a representation had been seen before, not whether meaning was preserved, and the authors present it as a first quantitative attempt with a deliberately small element set ("an example, rather than a definitive list"). The setting is US oncology; the elements are 12; results, units and categorical values were not assessed; neither FHIR nor openEHR appears anywhere in the study, which examined EHR-internal representations rather than exchange formats. **The score is therefore not a loss rate**, and readings of it as "30% to 80% of information is lost in exchange" overstate it (one such reading is recorded at E-12). Three authors contributed as employees of CancerLinQ LLC, disclosed in the article.

D-12: EURIDICE — Who Writes the Specifications the Implementing Acts Will Reference

§ 1.1 records that the EHDS Regulation is in force and that its technical content is being concretised through Implementing Acts by March 2027. That statement is incomplete on its own, because an Implementing Act does not invent technical content. It references specifications produced elsewhere. This appendix records where they are produced, on what basis that can be asserted, and how far the observation carries.

What EURIDICE is. EURIDICE, European Interoperability Specifications for Digital Solutions in Healthcare, is described by both publishing bodies as a joint HL7 Europe and IHE-Europe initiative. Its stated focus is "creating well-architected, implementable and testable specifications offered to the European Commission for being referenced in the EHDS Implementing Acts". The stated method is to combine HL7's FHIR expertise with the IHE methodology and to align with EU-funded projects, national priorities and standards developing organisations.

The specification catalogue, as published.

SPECIFICATION	SUBJECT
EU Patient Summary	The cross-border summary record
EU Laboratory Report	Laboratory results
EU Hospital Discharge Report	Discharge documentation

SPECIFICATION	SUBJECT
EU Imaging Report	Radiology and other imaging reports
Manifest-Based Access to DICOM Objects (MADO)	Access to the imaging objects themselves
EU Medication Prescription and Dispense	Prescription and dispense
EU Health Data API	The access interface

All are FHIR implementation guides. Two of them, the EU Imaging Report IG (STU 1 ballot) and the EU Health Data API IG, were placed in coordinated ballot by HL7 Europe and IHE-Europe with a deadline of 30 April 2026. Voting was restricted to HL7 Europe affiliates and IHE-Europe National Deployment Committees, and comments were open to anyone.

What it carries. It converts the regulatory finding of § 1.1 from a declaration into a mechanism. The specifications an Implementing Act is expected to reference are being authored as FHIR implementation guides, by the two organisations that hold the European profiling and testing machinery, on a timetable that runs ahead of the March 2027 deadline.

That is the concrete backing for the floor asymmetry of § 1.6.1: an openEHR-held record must cross a model boundary for European exchange, a FHIR-held record need not. It also backs the judgement that a change in the exchange requirement is not a near-term prospect.

What it does not carry. It says nothing about clinical quality, modelling depth or persistence, and it is not an argument against openEHR. It addresses the exchange layer, which this document treats as settled in any case and which is not the contested question. Nor is a ballot a legal act: the Commission has not adopted these guides, and the appearance of a specification in the EURIDICE catalogue is a proposal to the Commission, not a decision by it. Statements about what will in fact be referenced in 2027 remain forecasts.

Evidence grade B. The propositions rest on published artefacts of the two bodies, the initiative's own specification catalogue and the coordinated ballot announcement, read for what they state. They are independently inspectable but not constitutive: only the adopted Implementing Act would be constitutive, and it does not yet exist. The grade attaches to the proposition that these specifications are being produced by these bodies for this purpose, not to any proposition about the outcome.

Interests, disclosed at the point of use. The author of this assessment is CEO of HL7 Deutschland and a Senior Expert and Community Event Manager at HL7 Europe. One of the two bodies producing EURIDICE is therefore an organisation in which he holds a role, and the finding recorded here supports a document that recommends a FHIR-native default.

The mitigating consideration is the character of the proposition rather than the standing of the source. What is asserted is that named specifications exist, are published as FHIR implementation guides and were balloted on a stated date. That is checkable by inspection against the artefacts, and independent of who reports it. Any reading beyond that, that the Commission will adopt them or that this settles the storage question, is not asserted here, and would not be carried by these sources in any case.

Correction recorded. Up to Release 1.5.0 this document set named EURIDICE at one place only, in § 1.1, as an example of complementary architectures and methodologies for interoperability levels 4 to 6 and as neutral with respect to the choice of standard. Both halves of that characterisation are wrong against the sources: EURIDICE produces FHIR exchange specifications, which sit at the levels both standards already serve, and it is not standard-neutral. The sentence has been corrected and the description moved to § 1.1 in its own right.

D-13: AKTIN — Register Reporting Built on Care Documentation

Why it is here. § 2.6.1 records that form-driven capture is often asked for derived, purpose-specific facts that no ward documents in that form, and that deriving them is itself a mapping. The emergency-department network AKTIN publishes an implementation guide that builds register reporting on that observation instead of working around it, which makes it the worked example for the point.

What the guide describes. Questionnaires and QuestionnaireResponses are used as the reporting mechanism and are described as complementing, not replacing, classical profiling: they allow summary, register-specific and rapidly adaptable collection, with pre-population from data already held. The record the pre-population draws on, the AKTIN Record, is assembled from ordinary clinical building blocks such as `Observation`, `ServiceRequest` and `MedicationAdministration`, and the guide states the governing principle as no shadow data management separate from care.

Register-specific derivations are treated as computations over that record rather than as separate documentation: the guide's own example is age. One register requests it in completed days, another in months and a third in categories, all derived from the source dates. Items are annotated with published terminologies, LOINC for observations and SNOMED CT and ATC for clinical concepts, in preference to closed answer lists, which is the semantic-annotation discipline described at § 2.5.2 applied to forms.

What it carries and what it does not. It carries the proposition that register reporting can be organised as a projection of care documentation, with the derivation named and specified rather than left to a local convention. In such an arrangement a Questionnaire is a question over a record rather than a substitute for one.

Evidence grade B for what the guide specifies, as a published implementation guide verifiable by inspection. It carries no proposition about either standard's suitability, no measurement of how much the derivation costs, and no evidence about the quality of what the registers receive.

Interests. The guide is published through the ART-DECOR publication environment, and the author of this assessment is a member of the ART-DECOR Expert Group. The propositions recorded here are readings of the published guide and checkable against it.

D-14: Conformance Without Interoperability, Measured Between Two Implementation Guides (Kramer & Moesel 2023)

Why this appendix exists. The proposition that conformance is not interoperability is stated at § 1.5 and illustrated by the German corpus. This appendix records a peer-reviewed method applied to two widely used FHIR implementation guides, US Core and the International Patient Summary, showing that the issue is not uniquely German.

What it is. Kramer MA, Moesel C (2023): *Interoperability with multiple Fast Healthcare Interoperability Resources (FHIR) profiles and versions.* JAMIA Open 6(1), ooad001. Both authors are at the MITRE Corporation, which builds FHIR tooling for United States federal programmes and maintains two of the validators the paper names. No competing interest declared, no external funding. The article is a methods contribution, not a review: the authors propose an analytic instrument, the **FHIR Interoperability Table**, and demonstrate it on one exchange.

The instrument, in the terms the assessment needs. Two concepts carry it.

- A profile's **catchment** is the type or category of data the profile is *intended* to represent, read from the profile's narrative, because FHIR provides no computable expression of it.
- A profile's **admittance** is the set of resources that actually validate against it, objectively determinable from the `StructureDefinition`.

Admittance may be wider than catchment but must never be narrower, or the profile excludes what it exists to carry. Two profiles then stand in one of five relationships: equivalent, subset, superset, overlapping, disjoint. An exchange from a sender enforcing A to a receiver enforcing B is **feasible** where the relationship is equivalent or subset, **potentially infeasible** where it is superset or overlapping, and **infeasible** where it is disjoint. The judgement is made element by element, and one infeasible element makes the resource infeasible.

The finding that bears directly on § 1.5. The authors state in the methods section, as a premise rather than a conclusion, that confirming two systems conform to their respective specifications does not imply that the two systems can interoperate. That is the proposition of § 1.5 in the words of an independent team, in a peer-reviewed venue, and supported by a worked case rather than by argument.

The worked case. US Core 5.0.1 as sender against the International Patient Summary 1.0.0 as receiver: the American national guide against the international summary guide, both conformant FHIR, both mature, both widely implemented. The full analysis would cover 23 resource types, 43 US Core profiles and 26 IPS profiles; the authors report three resource types, four profile pairs:

RESOURCE	SENDER PROFILE	RECEIVER PROFILE	RESULT	WHY
Patient	US Core Patient	IPS Patient	Potentially infeasible	`birthDate` required in IPS, not in US Core; `contact.relationship` bound to different value sets with a `required` binding on the IPS side; `communication.language` extensible in US Core against required in IPS
Condition	US Core Condition Encounter Diagnosis	IPS Condition	Potentially infeasible	`clinicalStatus` required in IPS, optional in US Core
Condition	US Core Condition Problems and Health Concerns	IPS Condition	Potentially infeasible	as above; also differing data-type restrictions on `onset[x]` and `abatement[x]`, and extensible category bindings drawn from different code systems
MedicationRequest	US Core MedicationRequest	IPS MedicationStatement	**Infeasible**	Different base resource. IPS represents reported medication as `MedicationStatement`, and does not require support for `MedicationRequest` at all

All three resource types examined fail, one of them outright. The authors also show the converse case, whether one sender could satisfy both guides simultaneously. A common Patient profile is achievable at the price of losing terms present in US Core and absent from IPS. Simultaneous support of patient medications is not achievable at all.

Two further statements this assessment uses.

- **On binding strength, from the tooling side.** The authors record that `preferred` and `example` do not constrain potential values, that `extensible` is typically treated by algorithms as if it did not constrain, and that **only a `required` binding definitively limits the admittance of a profile**. Section 2.5.2 argues the same point from the specification; this is the same point from what validators actually do, and it is the sharper of the two because it describes behaviour rather than intent.
- **On what an implementation guide is worth.** Their closing observation is that infeasibility does not arise unless `StructureDefinition`s are enforced, that exchange is easier where they are not, at the price of participants having to process resources of any kind, and that where compliance is not expected or required, an implementation guide adds little. That is the conformity-assessment argument of § 1.5 and Document 5 from the opposite direction.

Evidence grade. A for the relationships computed between the two named profile versions, which are direct readings of published artefacts and reproducible by anyone with the same two guides and the FHIR validator. **B** for the method as a general instrument, which is proposed and demonstrated but not yet evaluated by third parties.

No grade for any proposition about how often such conflicts occur in the field. The paper analyses three resource types of twenty-three, chosen for brevity and not by any sampling rule. It therefore establishes that the conflicts exist and are systematic in kind, not their prevalence.

What is expressly not taken from it. Not a claim that US Core and IPS cannot exchange data in practice. Systems bridge these gaps routinely with transformation, negotiated subsets and relaxed validation, and the authors say as much. What the paper establishes is that the bridging is work that the conformance of both sides does not remove, which is the entire content of § 1.5. It is also worth recording against the interest of this assessment's own argument: the finding applies to the recommended standard, in its most mature guides, and it is not softened here for that reason.

D-15: openEHR Template versus FHIR Questionnaire/SDC — The Derivation Behind § 2.6.1 of Document 3

This appendix carries the derivation whose conclusions stand at § 2.6.1 of Document 3. It was moved here on 7 September 2026 to shorten Document 3; the text is that of the section as it stood, with the cross-references adjusted to this document and nothing else changed. The objection it answers reads: *openEHR covers FHIR's Questionnaire/QuestionnaireResponse equivalents only.* The verdict is: **accurate about the form, wrong about the model.**

What SDC is, since the objection rests on it. Structured Data Capture is an HL7 implementation guide that extends `Questionnaire` and `QuestionnaireResponse` for form-driven capture. It adds advanced rendering, pre-population from existing data, calculated expressions, conditional display, modular reuse of form fragments, and the extraction of answers into clinical resources by way of `StructureMap`. It is a set of extensions and conformance rules on two resources, not a separate standard (glossary at Appendix F of Document 3; the guide is listed in Appendix G below).

The objection compares two different layers, and that is why it half holds. openEHR carries a reference model, archetypes that constrain it, and templates that compose archetypes for a use case. FHIR carries resources, profiles and logical models that constrain them, and Questionnaires that define forms.

The **form** layers correspond: an openEHR template and a FHIR Questionnaire both state what is asked, in what structure, with what value sets, and both produce an instance. The **model** layers correspond as well: an archetype and a profile or logical model both constrain a general model to a clinical concept.

The objection sets openEHR's model layer against FHIR's form layer. Read at the form layer it is essentially right, and the assessment says so. 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 alongside a FHIR artefact in the Symptoms Implementation Guide, the counterpart carried beside it is the Logical Model, not a Questionnaire (E-7 below).

Put as a question, the objection runs: if capture depth is all that matters, FHIR Questionnaires with SDC deliver the same. They offer nested groups, terminology binding, conditional logic (`enableWhen`), modular reuse, calculated and repeatable structures. For pure *depth of capture* that is accurate, and the argument "FHIR alone cannot represent this depth" is not tenable in that generality. Three differences survive the comparison. They are narrower and more concrete than the debate usually makes them, and they are taken one at a time below.

First: persistence target and query capability. Take a graded finding recorded on an intake form, the functional class of a patient with heart failure. Both routes capture it, and both persist it. What differs is what can then be asked of the store with the standard's own query mechanism.

In openEHR the grade is a node in an archetype, and AQL selects it by its archetype path. The same query returns the grade from every template that embeds that archetype, because the path belongs to the shared model rather than to the form. In FHIR the grade sits in a `QuestionnaireResponse` item identified by a `linkId`. R4 defines no search parameter for item answers or for `linkId`, and R5 adds only `item-subject`. In the two releases the European and German profiles are pinned to, a server therefore cannot be asked for the grade directly.

Three routes lead out of that, and each is worth naming. SDC extraction writes the answer into an `Observation` by way of `StructureMap`, after which it is queryable as an ordinary resource. A projection layer such as SQL on FHIR renders the item as a column for analysis. And the current specification build adds composite search parameters that pair `linkId` with the answer value, which closes the gap at the source once servers implement it.

Two things follow. The FHIR route reaches the same place by one further step, and that step is itself a mapping with the obligations of § 1.6 of Document 3. And the openEHR query key is shared across forms, while the FHIR key stays local to one form until those composite parameters are released and implemented.

One bound belongs on that advantage, and it is about products rather than about the language. AQL is portable by design, but a query is only portable to a store that speaks it. In the vendor survey of eleven openEHR repositories, six declare AQL, four are queried in SQL and two through a formalism of the vendor's own, one product declaring both AQL and SQL (Delussu et al. 2024; A-2 above). The comparison above therefore holds between the specifications as written. Between the products as built it holds for the majority of openEHR repositories rather than for all of them.

Second: where the content models live, how far they are shared, and what is countable. The frequent claim that FHIR Questionnaires lack a governance procedure does not hold. A Questionnaire can be authored by a responsible body, reviewed by a working group, bound to published value sets, packaged in an implementation guide, balloted, versioned and published through the official registries. National programmes do exactly that. The difference is not the presence of machinery. It is the layer at which reuse is organised, and how far it can be verified from the artefacts.

On the profile layer, reuse is machine-readable. Every `StructureDefinition` names its `baseDefinition`, so the derivation graph of a corpus can be computed. It shows how many profiles constrain a national base profile rather than an HL7 base resource, how deep the chains run, and which base profiles carry the ecosystem. That is what a profile hierarchy is, and the German corpus of Document 5 consists of exactly the artefacts required for it (Document 5, D-10.7).

For the German corpus the count exists, in Document 5 (D-10.7), and its conclusion is enough here. Reuse is real, and it is local: most derivations stay inside one package, a few programme-internal hub artefacts carry a large share

of them, and the shared national layer is not the structural parent of the landscape. That is a finding about German governance rather than about FHIR.

On the openEHR side the equivalent is measurable in principle, because a template names the archetypes it composes, though no such count is published.

On the Questionnaire layer there is a mechanism but no library. SDC defines modular forms. A display item carries a `subQuestionnaire` extension naming another Questionnaire by canonical URL, version-specific as the guide recommends, and the `$assemble` operation expands the references recursively into one form. Composition is therefore specified. What the specification does not establish is a curated, shared holding of form modules. The guide itself goes no further than noting that research organisations keep question libraries, and ships one as an example artefact.

Whether such a library exists is consequently a matter of what a community builds beside the standard, and the counter-example is not hypothetical. The REDCap Shared Library curates data-collection instruments through an oversight committee that reviews them for relevance, correct functioning and copyright before release, and projects import them into their own forms. A form library is achievable. In the FHIR case it is not part of the standard, and in the openEHR case it is.

Third: what an instance needs in order to be read. A `QuestionnaireResponse` item carries a `linkId` whose meaning is defined by the Questionnaire it references, and the reference is a canonical URL that need not name a version. An openEHR composition carries at-codes whose meaning is defined by the archetype or operational template (Document 3, § 2.5.1). The dependency is therefore of the same kind on both sides, which is the honest starting point.

The difference is its reach. An at-code sits under an archetype path that is stable across every template embedding that archetype, so the same node keeps its identity across forms. A `linkId` is scoped to one Questionnaire, and two forms asking the same clinical question can and usually do use different ones. Neither is self-sufficient without its definition artefact. One is portable between forms and the other is not.

A fourth observation belongs to forms rather than to either standard: what forms are asked to capture. Form-driven capture is frequently used for purpose-specific, derived facts rather than for the clinical record as documented. A register asks "were antibiotics administered", a legitimate and answerable question that no ward documents in that form, so the answer has to be derived from medication administrations recorded for care. An age in completed days, months or categories is derived from a date of birth and an encounter date.

Deriving those answers is itself a mapping, with the obligations of § 1.6 of Document 3, and it appears in neither standard's specification because it is a property of the use case.

The German emergency-department network AKTIN builds its register reporting on exactly this arrangement and states the principle plainly (D-13 above):

- Registers draw their summaries from standardised clinical building blocks: `Observation`, `ServiceRequest`, `MedicationAdministration`.
- There is to be no shadow documentation beside care.
- Questionnaires are pre-populated from that record for the register-specific questions, and annotated with LOINC, SNOMED CT and ATC codes rather than with closed answer lists.

The point for this comparison is that **a Questionnaire is not automatically a form for capture at the point of care**. Where it carries derived content, the mapping it hides sits between the record and the question, not between two standards.

Empirical finding on modelling depth. The claim that FHIR is structurally too shallow for research-grade clinical modelling can be tested against a systematic body of reviews. It does not withstand the test.

Three systematic reviews of differing methodology carry the finding (in detail D-8 above):

- **Ayaz et al. 2021**, a review per the Kitchenham methodology.
- **Vorisek et al. 2022**, a PRISMA review prospectively registered in PROSPERO.
- **Tabari et al. 2024**, a scoping review per PRISMA-ScR.

Between them they document FHIR implementations in oncology, genomics, infectious disease and imaging (D-8 above). Further uses named in the debate, cancer survivorship, genotype–phenotype association studies and machine-readable guideline formalisation (EBMonFHIR, CPG-on-FHIR), are not among the domains those reviews record.

FHIR is not used predominantly for exchange in this body of work. It is used for data standardisation (41%), for **primary data capture** (29%), for cohort recruitment (14%) and for analysis (12%). In Tabari et al., 81% of the implementations examined are dynamic pipeline models rather than static mappings.

The conclusion is bounded, and it is enough for the argument here: depth of representation is not what separates the two standards, so an architecture decision cannot be made to rest on it.

What this demonstrates, and what it expressly does not. It does **not** demonstrate that FHIR is superior to openEHR in modelling depth. None of the three reviews includes openEHR as a comparator, and none measures effectiveness; Vorisek et al. expressly concede publication bias.

What it does is remove one argument from the debate. The claim that FHIR is fundamentally unable to represent the depth needed for secondary use does not hold as a blanket statement, and a tender that rests on it is resting on nothing. The differences that remain are the three named above: query key, location of the shared model layer, and the reach of the instance's pointer. None of them is about expressive power.

An honest delimitation. None of the three points is a FHIR deficit in principle. With logical models, consistent profiling, SDC extraction and a repository with SQL on FHIR, each can be answered inside the FHIR ecosystem. The pending composite search parameters answer the first one at the source. Nor is the openEHR side as settled as its description suggests: the curated library is a real asset, and how much of it is reused rather than re-modelled locally is unmeasured.

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 this means that 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.

Appendix E: Source Criticism of the 2024–2026 Materials

This appendix classifies the position papers, frameworks and industry contributions added in the course of consolidation by character, interest situation and robustness (E-1 to E-6). It documents two factual examinations of contested claims from the specialist debate (E-7, E-8), assesses one cross-standard systematic review (E-9), one intergovernmental policy paper (E-10), one statutory advisory position (E-11) and one ministry-commissioned national consensus study (E-12). It is deliberately kept separate from Appendix D, because the issue here is what the respective documents and claims can carry *as evidence*.

Every entry answers the same questions, and the grid below states the answers so that the appendix can be audited without being read end to end. The prose entries that follow give the reasoning; where the two differ, the prose is the fuller statement and this grid the index to it.

#	SOURCE	CLAIM EXAMINED	FINDING	GRADE	IMPLICATION FOR THE ASSESSMENT
E-1	Tsafnat et al. 2024, JMIR editorial	That the three standards occupy three distinct domains, openEHR persisting and FHIR exchanging	The division is asserted from design intent, not derived from evidence; three of five authors are affiliated to the camps whose roles they assign	D, editorial	The three-domain thesis is carried as a position to be examined, not as a finding (§ 2.4.1)
E-2	Siemens Healthineers white paper 2026	Vendor-side comparison and implementation recommendation	Self-declared not peer-reviewed; a supplier writing on architecture it also sells into	C/D	Read as market signal, not as evidence (§ 1.4)
E-3	ÖGTelemed position paper 2026	That vendor-neutral health data require an open persistence standard as an award criterion	The paper admits openEHR and FHIR alike for that role; it is from a specialist society, not a chamber or an authority, a distinction widely confused in the debate	C	Carried as an institutional position with its standing corrected (§ 1.4)
E-4	Australian Digital Health Agency framework 2026	National governance arrangements for digital-health standards	An official framework; strong for what it says about governance, silent on effectiveness	B for governance statements	Supports the governance argument, not the architecture choice
E-5	Fuhrmann et al. 2026, conference poster	An openEHR archetype for intraocular pressure as a vendor-neutral model	Evidences the CKM review process, not an architectural claim	D	Carried as evidence of openEHR's governance machinery (§ 1.4)
E-6	Devitt 2026, blog	That openEHR faces a "Red Hat problem" of tooling capture	Runs against the apparent interest of an openEHR-leaning author, which is recorded but does not raise the grade	D	Supplies a distinct lock-in mechanism (§ 1.3). The author's premise of openEHR's technical superiority is expressly not adopted

#	SOURCE	CLAIM EXAMINED	FINDING	GRADE	IMPLICATION FOR THE ASSESSMENT
E-7	HL7 Symptoms Implementation Guide and surrounding claim	That openEHR archetypes are displacing FHIR logical models	Refuted against the primary source: the archetype accompanies the logical model as an informative artefact	B, direct reading	Evidences convergence rather than displacement (§ 2.4.1, § 2.6.1)
E-8	Dublin convergence process and the Swiss joint working group	That convergence between the organisations lowers transformation cost	Convergence in publication is evidenced; a cost effect is not, and the Swiss case is an announced proof of concept	D for the Swiss initiative	Convergence announcements are not evidence of maturity or of lower cost (§ 1.5, § 1.6)
E-9	Marfoglia et al. 2026, systematic review	The distribution and adaptation of all three standards across 99 studies, and a hybrid-ecosystem conclusion	The counts hold over the review's own inclusion list; the hybrid conclusion is derived from design intent, not from an outcome or cost comparison	B for the counts, no grade for inference to deployment	Carried as a named counter-position with three qualifications (§ 1.2, § 1.5, § 2.7)
E-10	OECD 2026, interoperability working paper	The multi-dimensional policy framing, governance and implementation barriers, lifecycle mapping, and an illustrative FHIR/openEHR/OMOP allocation	Policy synthesis rather than an architecture comparison; one FHIR statistic overstates its cited study, and the economic estimates are heterogeneous projections	B for the framework and survey findings; C for third-party projections; no grade for the technical taxonomy	Used for the broad interoperability frame and governance findings; the standards allocation and cost conclusion are not adopted
E-11	Interop Council 2026, position paper	That a statute should prescribe semantically standardised data holding at the logical layer while leaving databases and architectures unregulated	The demand divides the term the consultation used whole; it reaches only bindingly determined data contents; openEHR and any persistence standard go unmentioned	A for the wording; B for the positions	Carried at § 1.2 and § 1.6.3 as a recommendation, not a determination
E-12	Piera-Jiménez et al. 2026, Delphi study	That openEHR is the most suitable standard for Spain's	A ministry-commissioned expert consensus,	A for the wording; B for the consensus	Carried as the weightiest institutional counter-position to this

#	SOURCE	CLAIM EXAMINED	FINDING	GRADE	IMPLICATION FOR THE ASSESSMENT
	(Lancet Digit Health)	national digital medical record, with FHIR at the exchange layer	not an outcome comparison, and by its own words not evidence-based; authorship overlaps with openEHR leadership, openly declared	figures; no grade for the suitability verdict	assessment's default (§ 1.1, § 1.4, § 2.4.1)

Limitations are stated inside each entry rather than in this grid, because they are the part that does not compress. An interest declared in a paper, a sample that was not drawn systematically and a technical horizon that has moved are three different weaknesses, and a single column would flatten them.

E-1: Tsafnat et al. (2024) — the Three-Domain Editorial

Tsafnat G, Dunscombe R, Gabriel D, Grieve G, Reich C (2024): *Converge or Collide? Making Sense of a Plethora of Open Data Standards in Health Care.* J Med Internet Res 26:e55779. DOI: 10.2196/55779.

Publication type: Editorial, expressly **not peer-reviewed**. *Evidence grade D: conceptual positioning, no empirical evidence.*

Thesis. openEHR for clinical primary documentation and administration, FHIR for exchange, OMOP for longitudinal analysis; the informatics community should converge on these three standards.

Substantive critique.

1. *Transformation costs systematically undervalued.* The thesis presupposes semantically low-loss transitions between three layers. The authors concede losses but confine the examples to technical cases. The actual problems remain unaddressed: governance of the transformation logic, quality assurance across system boundaries, and accountability for clinically relevant semantic losses.
2. *Patient-safety dimension absent.* Semantic degradation at the layer transition can have clinical consequences (§ 1.6.2). This dimension is entirely missing.
3. *Proprietary HIS inertia not addressed.* The model implicitly presupposes that primary data exist in one of the three open standards. The structural reality, dominance of proprietary HIS with limited export interfaces (§ 1.3, Dimension 3), is not discussed. The thesis thereby operates above the level at which the actual interoperability problem lies.
4. *Historical division of roles reproduced uncritically.* The assignment of FHIR to the pure exchange domain reflects a state of affairs that productive FHIR-native CDRs (Appendix D-7) have overtaken.

Formal critique (subordinate, but to be stated). The authorship consists of institutional leaders of the standards respectively recommended, among them the then CEO of openEHR International, the founder of HL7 FHIR and the CEO of an OMOP/OHDSI service provider. The conflicts of interest are declared in the article. The paper is thus to be read as a position paper by standards representatives, not as independent analysis.

Peer-reviewed counter-position. Kapitan D et al. (2025): *Data Interoperability in Context: The Importance of Open-Source Implementations When Choosing Open Standards.* J Med Internet Res (Viewpoint). The authors argue that open standards are a necessary but not sufficient condition for interoperability; the ecosystem of open-source implementations must be considered as well. The critique advanced here does not, therefore, stand alone.

Usability.

PURPOSE	USABILITY
Evidence that the complementarity thesis is a widespread position	Citable, with a note on the authors' interests
Empirical evidence for the suitability of the three standards	Not citable
Evidence of willingness to converge between the standards communities	Citable as an indication, not as proof
Reference foil for the substantive counter-presentation	Suitable

E-2: Siemens Healthineers — Technical White Paper (June 2026)

Schabetsberger T, Mangesius P (2026): *HL7 FHIR and openEHR: A Comparative Analysis, Implementation Strategies, and Recommendations for Healthcare Organizations and Regions.* Technical White Paper v1.0, Siemens Healthcare Diagnostics GmbH, Vienna.

Character. A vendor paper that makes its own status transparent: no peer review, no systematic literature appraisal, no independent editorial scrutiny. *Evidence grade C/D: substantively solid, but not independent.*

Core thesis. FHIR is the mandatory baseline — regulatorily mandated, enforced by the market and supported by practically every relevant EMR platform. openEHR is a valuable but *additional* layer of semantic depth for organisations with research-intensive or longitudinal data requirements. The recommendation: FHIR first, openEHR selectively.

Analytical value. The paper is currently the most detailed structured juxtaposition of the two standards and contains substantive contributions that this document takes up:

- A differentiated **three-stage taxonomy of openEHR integration**: native use of the openEHR data model (DIPS Arena), a standalone CDR *alongside* the primary system, and connection of conventional systems via a FHIR bridge. The middle stage is named as the pattern most widespread in Central Europe.
- A precise naming of the **structural reason for mapping losses**: openEHR archetypes are maximal models, FHIR resources deliberately minimal datasets for the most frequent use case. A semantic levelling ("flattening") in transformation is therefore inherent, not the result of poor implementation. That is the sharpest available formulation of the problem treated in § 1.6, and this document adopts it.
- A time estimate of **12 to 24 months for the clinical modelling phase** before productive operation of an openEHR architecture, consistent with the FIKS finding (Appendix D-5).

Critical classification. Two reservations must be made. First, the conclusion ("FHIR first") coincides with the publisher's market position; Siemens Healthineers is not a neutral third party on the question of which persistence architecture prevails.

Second — and more important for procurement practice — this paper is the **technical evidentiary basis of the ÖGTelemed position paper** (Appendix E-3), which cites it as a source at five points. A vendor document without peer review thereby carries central statements of a national position paper. That devalues neither the paper nor the position paper, but it must be disclosed in the weighting.

Usability. Robust and citable for the description of technical matters (integration stages, mapping mechanics, the market situation of EMR FHIR support). For normative conclusions on the choice of standard, only with a vendor caveat.

E-3: ÖGTelemed — Position Paper on Vendor-Neutral Health Data (June 2026)

Austrian Society for Telemedicine and eHealth (ÖGTelemed) (2026): *Positionspapier zu herstellerneutralen, applikationsunabhängigen Gesundheitsdaten.* Version 1.0, 24 June 2026, 17 pages.

Character. A position paper by a specialist society — **not** a professional chamber or public authority. *Evidence grade C: an institutional position with recommendatory character, not a regulatory act.* This distinction matters, because the paper is occasionally cited in the debate with greater weight than its institutional status carries.

Core demand. A vendor-neutral, application-independent persistence layer for clinical data; a national coordination and maintenance process for clinical information models; anchoring of open persistence standards as a **mandatory criterion in award procedures**, particularly in the context of upcoming SAP IS-H replacements and new HIS tenders.

Why the paper is relevant to this analysis, and in an unexpected direction. The paper comes from an environment one would assign to the persistence argument, and yet it supports a core thesis of this assessment (§ 2.4). It expressly records that **both openEHR and HL7 FHIR** can be used for the vendor-independent persistence of clinical data. It names the discussion of FHIR as a persistence standard within the HL7 community, and states that the fundamental goal, an application-independent data layer, is the same under both approaches.

The recommendation consistently reads "openEHR **and/or** FHIR", not "openEHR".

The actual dividing line the paper draws. Its strongest argument is an architectural one: a FHIR interface on a still-proprietary primary system does not solve the lock-in problem, because the data remain in the vendor-bound data model. FHIR as a mandatory exchange layer and vendor-neutral persistence do not exclude each other; they presuppose each other. That is precisely the limit of Option A (§ 1.6.1) and coincides with § 1.5 ("conformance is not interoperability").

Critical classification.

- The figure of the target picture (Table 1 of the paper) nevertheless reproduces the domain separation (persistence *and* exchange *and* OMOP) and thus stands in a certain tension with its own body text.
- The technical burden of proof rests predominantly on the Siemens white paper (Appendix E-2), cited at five points. Peer-reviewed evidence for the central cost statements is lacking.
- The selection of international references is one-sided: OneLondon, Catalonia, Slovenia, Norway and UKSH/HiGHmed are cited as openEHR successes; FHIR-native persistence instances (Wales, MII DIZ) are absent, although the paper conceptually acknowledges FHIR persistence. That is not an inaccuracy, but it is a gap in the evidentiary base.
- The cost statement on the openEHR modelling phase (12–24 months) stems from the Siemens paper, not from the society's own data collection.

E-4: Australian Digital Health Agency — National Framework for Digital Health Standards (May 2026)

Australian Digital Health Agency (2026): *National Framework for Digital Health Standards. Connecting Australian Healthcare — National Healthcare Interoperability Plan.* ISBN 978-1-7643527-2-7, May 2026.

Character. An official national framework by a government agency. *Evidence grade B for policy and governance statements; no empirical evidence on standard effectiveness.* It is currently the weightiest official source that lists openEHR as a recognised standard alongside FHIR.

What the document says.

- Both standards are recognised: openEHR is listed as a standards development organisation alongside HL7 Australia, SNOMED International and GS1; the framework notes an increasing focus on openEHR and the value of reflecting the clinical voice in data structures.
- The Australian clinical dataset **AUCDI** (Australian Clinical Data for Interoperability) is expressly described as underpinned by *both* standards, FHIR and openEHR archetypes.
- At the same time, the modernisation of My Health Record is described as **FHIR-native**, and the government priorities (Sparked FHIR Accelerator, AU Core, AU Patient Summary) are FHIR-driven throughout.
- In the standards taxonomy (section 5.2, Figure 2), FHIR is assigned to the category *Information Exchange*, openEHR together with OMOP to the category *Content*.

Analytical finding: the internal tension. The framework perpetuates the historical division of roles ("FHIR exchanges, openEHR models content") into an official policy document, invoking in footnote 8 the Tsafnat editorial (Appendix E-1) — that is, a non-peer-reviewed source with declared conflicts of interest. In the same document, however, the central national infrastructure is described as FHIR-native.

That is not a contradiction in the strict sense, but it is a telling pattern: **the domain separation survives in the standards taxonomy while implementation practice undercuts it.** Precisely this pattern is the subject of § 2.4 and § 2.4.1.

Further statements relevant to this analysis. The framework expressly names as central adoption obstacles: commercial incentives for vendors to use proprietary rather than national standards; the costs of modernising legacy systems; and the absence of a national governance system. That is an independent, official confirmation of two findings of this assessment: § 1.3, Dimension 3, market inertia behind proprietary HIS internals, and § 1.5, the choice of standard as a necessary and not a sufficient condition. The source is attributable to neither of the two standards communities.

E-5: Fuhrmann et al. Conference Poster on the IOP Archetype (2026)

Fuhrmann L, Rodrigues I, Schargus M, Dicke C, Ong A, McNicoll I, Bakke S, Richter C, Leslie H, for the *openEHR eye care collab* group (2026): *Developing a Vendor-Neutral, Context-Inclusive Clinical Information Model for Intraocular Pressure.* Conference poster P.178.

Character. A conference poster (available in full), no peer review, no accompanying full-text publication. Produced within the openEHR eye-care collaboration and authored by CKM editors and openEHR clinical modellers; the interest is in the process it demonstrates. *Evidence grade D.*

Content. International consensus review of an openEHR archetype for intraocular pressure via the CKM platform: 32 specialists from ophthalmology and clinical informatics from 11 countries, two formal review rounds (21 and 23 reviewers respectively), publicly accessible comments, formal v1 publication announced.

What it demonstrates, and what it does not. The poster is **not** evidence on the architecture question: it says nothing about persistence architectures, mapping costs or implementation success. It is, however, a concrete, documented demonstration of **how CKM governance works** (§ 1.4, § 2.6.1). That is precisely the property openEHR claims as a distinct strength over FHIR profiling, and which this document acknowledges as such.

A comparable internationally consented and publicly reviewed modelling process is not structurally provided for in the FHIR ecosystem. It could be replicated with Logical Models and profiling, but it is not part of the design.

Qualification. The existence of a functioning governance process is no proof that the models arising from it reach the point of care. Between a consented archetype and productively captured data lies the entire implementation burden that Appendix D-5 documents.

E-6: Devitt (2026) — "Is openEHR the Linux of clinical data standards?"

Devitt D (28 May 2026): *Is openEHR the Linux of clinical data standards?* darrendevitt.com.

Character. The professional blog of a single author with a practice background in clinical data integration; no peer review, no source references, no collected data. *Evidence grade D.*

Why the contribution nevertheless carries more than its evidence grade. The author is expressly well disposed towards openEHR and considers the standard the better clinical data model. His conclusion, however, runs against this preference: openEHR will not achieve broader adoption beyond the market already captured, and for most organisations FHIR is "good enough".

Statements that run against the evident preference of their originator carry greater weight in source criticism than partisan contributions in either direction, not because they are therefore true, but because the most obvious direction of bias is ruled out. The contribution is thus to be read as an *admission against interest*, not as advocacy.

The thesis. Technical superiority did not drive adoption. openEHR has succeeded where three conditions coincide: small or regional, state-run health systems, strong central coordination, and a long-term political willingness to fund clinical informatics. These conditions are preconditions, not accompanying circumstances. Where they fall away, the implementation fails. The major HIS vendors and hyperscalers, by contrast, are committed to FHIR; the US market sets the global investment direction and has decided regulatorily for FHIR.

The distinct analytical contribution: the "Red Hat problem". The contribution identifies one plausible path to broader openEHR adoption: vendors may encapsulate modelling complexity in tools, lowering the qualification threshold for implementation teams. If the resulting tools are proprietary, openness of the standard does not by itself ensure portability of the practical implementation path. This is a scenario and procurement risk, not a demonstrated market outcome.

This is a lock-in dynamic arising from the very solution of the skills problem — and it sharpens § 1.3, Dimension 1. Platform lock-in is usually described as *market concentration* (few vendors, vendor-survival risk). Devitt's point is structurally sharper: the more successfully a vendor lowers the entry hurdle, the stronger the dependency on precisely that vendor. The success of the standard and the openness of its ecosystem thus stand in a potential goal conflict.

For procurers a concrete test question follows: where an openEHR offering advertises reduced complexity and lower skill requirements, examine at what price that reduction is bought and how portable the results remain without the vendor's tools. That corresponds to the Gartner caution (Appendix D-6) to examine platform vendors for "genuine openness" and to regulate contractual data sovereignty at the end of the contract. Devitt supplies the reason why that caution is substantive, not formal.

The author defuses his own analogy, and that is its strength. For Linux, Red Hat solved the problem: a commercial vendor that encapsulated complexity, carried by an enterprise market that was large enough, and validated in the end by an acquisition in the tens of billions. For openEHR, nothing comparable is in sight: the market is smaller, commercial returns are harder to project, an acquirer of that magnitude not discernible.

The analogy thus ends precisely not in the success story its use in the debate usually suggests. Whoever uses the Linux analogy as an argument of confidence ("open standards prevail in the end") should take note that its most prominent proponent reads it exactly the other way round.

Convergence with the findings of this assessment. The contribution supports three findings, from an openEHR-leaning position and independently of the materials evaluated here:

- **§ 1.3, Dimension 3 (market inertia as the actual problem):** The dominant HIS vendors have no incentive to switch, and the regulatory environment of their largest markets has created none.

- **Appendix C-4, point 2 (skills shortage):** The dual qualification demanded corresponds to no established career path; projects without this expertise implement openEHR poorly and pay the full complexity bill without realising the clinical benefit that would have justified the choice. That is the same observation that Christensen & Ellingsen (Appendix D-5) demonstrate empirically.
- **Appendix C-4, point 3 (maturity of the commercial ecosystem):** The market leader in the openEHR platform segment is a company of around 200 employees; a comparable ecosystem of enterprise servers, cloud platforms, SDKs and regular connectathons exists on the FHIR side, and not on the openEHR side.

An independent data point on the classification of Wales. In May 2026, Devitt lists as real openEHR adoptions: Norway, Slovenia, Catalonia, Ireland (with a caveat of his own) and parts of Australia. **Wales is absent.** That an openEHR-leaning author does not count Wales among the reference implementations supports the classification made in § 1.1 from an independent direction.

The reading of the Dublin meeting. Devitt interprets the HL7/openEHR working groups that convened in Dublin in May 2026 as a departure from the long-held, partly unspoken position that technical superiority would bring adoption about by itself: the way forward runs *alongside* FHIR, not *around* FHIR. This interpretation is plausible but interpretative: it must be flagged as one observer's reading, not as the position of the organisations involved.

What is not adopted. The contribution's premise — that openEHR is technically superior to FHIR as a clinical data model, and that this dispute is settled — is **not** adopted here. It is asserted in the contribution, not substantiated, and it contradicts the analysis in § 2.4 and § 2.6.1. The robust difference lies in modelling philosophy (maximal set versus the 80-percent rule), persistence–query identity and clinical model governance, not in a superiority that could be established across standards.

Devitt's market diagnosis remains valid even if one does not share his premise; it is logically independent of it. Equally speculative, and not used here, is his closing suggestion that AI-supported tooling could lower the qualification hurdle and thereby change the outcome. He himself concedes that this would sooner accelerate than solve the Red Hat problem, unless the tools were themselves open.

E-7: Are openEHR Archetypes Displacing FHIR Logical Models? A Tested and Refuted Fear

In the FHIR-leaning specialist discussion the concern is voiced that HL7 has in effect declared openEHR archetypes the official route for Domain Analysis Models (DAMs) and will thereby displace its own FHIR Logical Models. The fear is testable, and the primary artefact refutes it in the form claimed.

The touchstone. The *FHIR R4 Symptoms Implementation Guide* (HL7 International / Clinical Interoperability Council, STU 1, build April 2026) is the most concrete case to date. In it HL7 places a Logical Model, a FHIR profile **and** an openEHR archetype side by side in one guide. The guide regulates the order of precedence explicitly:

- On the **Logical Model**, the guidance is that it is to be used as the authoritative reference for which symptom data should be captured; the FHIR profile serves to implement these concepts in systems.
- On the **openEHR archetype**, by contrast, it states that it is included in this guide as an informative document and as an additional source of clinical guidance. The standards status of the page in question is explicitly *Informative*.

Finding. Archetypes are migrating into HL7-published implementation guides. That is real, new, and a tangible result of the convergence efforts. But they do **not** replace the FHIR Logical Models: the Logical Model retains the authoritative status, and the archetype is carried along as informative additional clinical guidance. The displacement thesis is thereby refuted for this case.

Qualification. A broader HL7 governance decision not found here cannot be ruled out with final certainty. The concrete evidence speaks against it. Should such a decision exist, it would need to be evidenced. The burden of proof lies with the claim.

Why this section is here. This assessment tests claims regardless of which camp they come from. The fear tested here comes from the FHIR-leaning environment and proves, against the primary source, to be overshooting. The real core, that openEHR models are advancing institutionally into the HL7 space, remains, and is relevant to § 1.2 (point 5, standards-development process and institutional standing) and § 2.4.1. It evidences convergence, not displacement.

A constructive corollary: the gap is tooling, not capability. A related observation from the same FHIR-leaning environment reframes the concern productively. On this reading, FHIR Logical Models are functionally adequate for domain analysis; what they lack relative to openEHR archetypes is not modelling depth but tooling maturity, in two specific respects.

First, Logical Models have no graphical rendering in common use, although the FHIR publication toolchain already contains rendering logic for other artefacts that could in principle be applied to them. Second, there is no diagram-based authoring and no drawing tool, so that a converter from a general diagramming or UML notation into the Logical Model format would close the most visible gap.

Read this way, part of the relative attractiveness of archetypes rests on presentation and tooling rather than on substance, and the remedy lies within the FHIR ecosystem. Closing the tooling gap is an open task for the FHIR community, not evidence that Logical Models are the weaker instrument.

This is an informal community position (evidence grade D). It is recorded here as the constructive counterpart to the displacement concern above, and it is consistent with the finding of § 2.6.1 that the FHIR-versus-openEHR difference is one of maturity and tooling rather than of fundamental feasibility.

Tested: the rendering half does not hold in the form claimed. The corollary reports that Logical Models have no graphical rendering in common use. Measured against a published guide, that statement is too strong.

In the EHDS *Logical Information Models* guide (Xt-EHR), every model is rendered by the ordinary FHIR publication toolchain as a hierarchical structure table, in snapshot, differential and key-elements views. The table carries element names, cardinalities, data types, bindings and obligations, supplemented by a detailed per-element view and the machine-readable serialisations. That is a generated, uniform rendering of the model itself, and it is produced in the very domain this assessment is about.

What the same guide shows about the gap that remains. The guide does contain UML class diagrams with multiplicities, but at the level of the use case rather than of the individual model, and they are authored alongside the models rather than generated from them.

That is precisely the constellation the corollary's second half describes, and it survives the test intact. A diagram maintained separately from the model it depicts is a second artefact that can drift from it. There is still no diagram-based **authoring** of Logical Models, and no converter from a general diagramming or UML notation into the Logical Model format. The accurate formulation is therefore narrower than the corollary and narrower than its refutation alike: Logical Models are rendered, as tables; they are not authored, and not depicted, as diagrams generated from the model.

A further qualification: part of the authoring half is closed as well, from an unexpected direction. The corollary implies that no tooling produces FHIR Logical Models from clinically authored domain models. That is not accurate either.

ART-DECOR, the open specification environment used across the HL7 CDA and FHIR communities (Appendix F-7), maintains technology-independent **datasets**: hierarchical collections of clinical concepts with data types, value domains and terminology bindings. It also maintains use-case-specific **transactions** derived from them, and serves both through its version-aware FHIR endpoints as `StructureDefinition` logical models.

The distinction between the two is the analytically relevant part: a DECOR dataset carries no cardinality or conformance information and is therefore exported with `0..*`, whereas a transaction carries that information and transfers it into the logical model. The constrained, implementable logical model is thus a product of the use-case model, not of the concept glossary.

What this closes and what it does not. It does not deliver what the corollary literally asks for. There is no drawing canvas, no converter from UML or a general diagramming notation, and no rendering or import of logical models authored elsewhere. The path runs one way, from the governed model into the FHIR artefact. What it does deliver is the conversion function itself, from a different source notation. The residual gap is therefore narrow: diagram-based authoring for those who want to model by drawing, and a route back from artefacts authored outside the tool.

Source character and interest. The rendering finding rests on a published implementation guide and is verifiable by inspection — **evidence grade B**. The ART-DECOR finding rests on the tool's own documentation (docs.art-decor.org, FHIR server endpoints), with no independent evaluation — **evidence grade D**, the same grade this document applies to openFHIR in § 1.6.

The author of this assessment is involved in the ART-DECOR Expert Group. Both observations are recorded because they correct factual statements in the corollary, and they are graded accordingly. No recommendation in this document rests on either.

E-8: The Dublin Convergence Process — What Is Evidenced and What Is Not

The second joint annual meeting of openEHR International and HL7 International ("Converge & Collaborate") took place on 12 and 13 May 2026 in Dublin, following the first meeting in Amsterdam in 2025. The joint press statement of the two organisations reports substantial progress by the joint working groups established in Amsterdam on the technical alignment of the standards, and announces an extension to further standards partners and a permanent international forum.

Source character. The press statement is a **joint institutional self-presentation of both participating standards organisations**. It is reliable for the fact *that* the meeting took place, which working groups exist and what intentions were declared. For assessing whether the reported "substantial progress" actually occurred, it is not. No participant would have an interest in a contrary account. An independent assessment of the convergence process is not yet available.

What can be concretely demonstrated. The primary artefact identified here is the HL7 FHIR Symptoms Implementation Guide, in which an openEHR archetype appears as an *informative* companion to the authoritative FHIR Logical Model. This evidences convergence at the level of model publication. It does not demonstrate lower transformation cost (§ 1.6).

Convergence at national level. The pattern is now visible below the international forum as well. In mid-2026 openEHR Switzerland and HL7 Switzerland established a joint working group. Its task is to turn a vaccination showcase from the Swiss GovTech Hackathon into a reusable implementation guide, a "blueprint" for pairing openEHR persistence with CH-VACD/FHIR exchange. The transformation runs bidirectionally, using the FHIRConnect mapping specification and its openFHIR engine (openEHR announcement, "From proof of concept to a reusable blueprint," 2026).

This documents cooperation at national level. It remains an announced proof of concept rather than a productive deployment or outcome study (evidence grade D). Its bidirectional transformation is the cross-standard boundary found in Options D or E, depending on where the authoritative clinical content is captured; it is not Option B, which persists the additional representation in FHIR. Long-term openEHR persistence also presupposes continued availability and interpretability of the applicable archetypes and operational templates (§ 2.5.1).

Consequence for procurers. Announcements of standards convergence should not be read as statements about technical maturity or about reduced mapping costs. Convergence at the level of the standards organisations is no substitute for a conformance regime and for conformance testing (§ 1.5). Whoever justifies an architecture decision with "the standards are converging anyway" relies on a declaration of intent, not on a tested property.

What is deliberately \not\ claimed here. In the specialist discussion a counter-narrative circulates that the convergence process yields little substance. This assessment does not adopt it, because no citable source exists for it. The only publicly published observer from the FHIR-leaning environment who has commented on Dublin (Devitt 2026, Appendix E-6) reads the meeting, on the contrary, as significant. He interprets it as a pragmatic departure by the openEHR side from the expectation that technical superiority would bring adoption about by itself.

Both readings, official optimism and informal scepticism, are at present equally unevidenced. What alone can be recorded is the observation that the convergence process is also read **tactically** by those involved, and not exclusively on the merits; that is a data point for § 1.7, but not a finding about the technical substance.

E-9: A Systematic Review of All Three Standards Side by Side (Marfoglia et al. 2026)

What it is. A systematic review in the *Journal of Biomedical Informatics*, registered in PROSPERO (CRD42024623398), reported against PRISMA, open access under CC BY. From 2,102 records across five databases it retains **99 peer-reviewed studies published between 2021 and 2024** that apply FHIR, OMOP-CDM or openEHR in a healthcare or research setting. Each is classified by standard, scale, purpose, application domain, ETL tooling and coverage. It is the first review in this material that takes all three standards together rather than one of them at a time.

Source character. A peer-reviewed secondary study by an academic team with no declared standards affiliation, funded from an Italian national research programme. The authors disclose the use of a large language model for proofreading. **Evidence grade B for its counts**, which are propositions about a defined body of literature and reproducible from the published inclusion list; **no grade** for the extrapolation from that literature to deployment, which the review itself declines to make.

The counts that bear on this assessment.

- Standard distribution: **56 studies OMOP-CDM, 39 FHIR, 8 openEHR** of 99; the sum exceeds 99 because two studies use two standards.
- Setting: **86 of 99 (87%) are research**, 13 are clinical; **90 of 99 (91%)** apply the standard to a secondary purpose rather than to direct care.
- Purpose: data reuse 46, clinical decision support 23, data exchange 23, vocabulary definition 12, EHR design 6.
- Adaptation: 77 studies state whether they used the standard as-is or adapted it. Of these, **23 of 27 FHIR implementations (85%) adapted the standard**; **25 of 44 OMOP-CDM studies (57%) used it unmodified**; openEHR splits 4 to 4.
- Coverage: only **27 of 99 (27%)** report how much of the source dataset the standard actually covered. The authors state that the mapped share is so rarely reported for FHIR and openEHR that the standards cannot be compared on it.
- ETL transparency: **15 of 99 (15%)** describe their ETL tooling comprehensively, 37 partially.

- **None of the eight studies using openEHR reports any limitation.** The reviewers record this as an observation, without comment.

What the review concludes, and why it is carried here as a counter-position. Its closing recommendation is a **hybrid ecosystem** — openEHR to persist, FHIR to exchange, OMOP-CDM to analyse, which is the arrangement Document 3 advises against for genuinely new infrastructure.

The position is handled at § 1.2 rather than set aside. Three qualifications travel with it and are stated there. The role assignment is derived from design intent and not from any outcome or cost comparison, none being present in the review. The material is the peer-reviewed literature, from which commercially deployed systems are absent by construction, as the authors say. And the review's own coverage finding means that the transformation loss between the layers of the proposed ecosystem is unmeasured on its evidence base.

What it confirms that this assessment already carries. That no comparative study of outcomes or lifecycle cost exists, and the authors ask for one in their closing paragraph, which is Desideratum 21 in the words of an independent team.

That the openEHR literature is thin and self-reporting, the absence of any reported limitation across eight studies being the sharper form of what Appendix E already says about camp self-presentation. And that adaptation of FHIR is the normal case rather than the exception, which the German corpus shows from the artefact side and this review shows from the publication side.

What is expressly not taken from it. No statement about market share, installed base or deployment maturity. The counts describe publications. The review is also bounded to 2021 to 2024 and to English-language peer-reviewed work, and its authors name publication bias toward research-documented initiatives as a limitation of their own study.

E-10: OECD (2026) — Interoperability as Governance, and the Limits of Its Standards Taxonomy

What it is. *Interoperability in Healthcare: Towards an Interconnected Future* is OECD Health Working Paper No. 197, published on 24 July 2026. It combines desk research with semi-structured interviews conducted principally in 2025, consultation with standards organisations, a workshop and 22 country survey responses.

The report examines interoperability through people, policy, process and technology rather than comparing repository architectures. **Evidence grade B for what the OECD framework and survey report as policy findings; C for unreplicated quantitative estimates reported from third-party studies; no grade for the technical correctness of its illustrative allocation of standards.**

What reinforces this assessment. The OECD distinguishes technical, semantic, legal, regulatory and organisational dimensions of interoperability and treats governance, trust, adoption, workforce capability and data quality as necessary conditions. It reports inconsistent data-sharing practices, technical fragmentation and fragmented institutional coordination among the most frequently named barriers. It separately identifies uneven standards uptake, local extensions, vendor-specific endpoints, locked data models, incomplete content profiles and insufficient conformance testing.

This independently supports the scope statement in Document 3 and the proposition of § 1.5: choosing an open standard is one foundation for interoperability, not its completion.

The report also describes mapping as a recurring activity along the health-data lifecycle. Differences in source schemas and missing context affect later interpretation, and each connection creates work to make data fit for purpose. This supports the general boundary argument of § 1.6. Its stronger statement that mapping activities reduce granularity and quality is not adopted. A transformation creates a risk of loss; whether relevant meaning is actually lost depends on the source, target, mapping, use case, validation and governance.

The three-standard allocation, and what it establishes. The OECD's illustrative example presents FHIR as the structure and mechanism for exchange and integration with clinical repositories, openEHR as a specification for structured longitudinal records, and OMOP as the common data model for secondary analysis. Its later cross-border recommendation groups FHIR, IHE profiles and openEHR as open foundations and OMOP as an example for aggregated data.

This is evidence that the three-domain thesis has entered international policy discourse. It is not comparative evidence for the thesis. The figure is expressly illustrative and non-prescriptive; the openEHR role is sourced to openEHR International, while the OMOP role in the example is sourced to an industry blog. No outcome, lifecycle-cost or semantic-fidelity comparison between the proposed paths is supplied.

A material misreading of Kramer. Paragraph 20 states that FHIR's intended 80/20 rule has become 65/35, and then characterises the 35% as data elements requiring implementation-specific extensions outside the standard that cannot be fully shared between FHIR systems. **The source is Kramer's profile study, not Kramer and Moesel's compatibility study**, and the two are separate works.

Kramer MA (2023), *Reducing FHIR "Profiliferation": A Data-Driven Approach*, reviews 2,897 profiles from 125 implementation guides and concludes that FHIR's intended 80/20 rule is closer to 65/35. Base FHIR supplies approximately two thirds of the elements the guides studied needed, and profiling is used to define 33 to 38% of new data elements.

That is what the paper establishes, and it is not the OECD's proposition. Profiling actions include cardinality changes, datatype constraints, value-set bindings, fixed values, slicing and extensions. The paper separately reports that most extensions encountered were custom.

It does not show that 35% of all exchanged data elements are custom extensions, nor that all such elements are unshareable. The OECD sentence conflates dependence on profiling, use of extensions, reuse and actual cross-system compatibility. It is not carried into this assessment. Kramer and Moesel's direct compatibility analysis, a different paper by one of the same authors, remains the relevant source for the narrower conformance-without-interoperability claim (Appendix D-14).

Economic claims. The report collates projected annual benefits equivalent to 2.70–6.64% of national health expenditure and concludes that interoperability will finance itself over time. The underlying estimates come from different countries, baselines, periods and bundles of assumed benefit, including administration, research, artificial intelligence, personalised medicine and avoided harm.

They do not establish a general return on investment and say nothing about the relative economics of FHIR, openEHR or OMOP architectures. One figure is directly relevant as context: a Finnish projection places investment in harmonised data models based on openEHR and OMOP at up to EUR 2 billion over three to five years.

This demonstrates the order of magnitude assigned by that programme's analysis, not the cost attributable to either model and not a comparison with an equivalent FHIR architecture. The primary Finnish source would require separate appraisal before the figure could carry an economic finding.

What is taken from the report. Its multi-dimensional definition of interoperability; its policy-level finding that standards adoption must be accompanied by governance, compatible profiles, terminology, conformance testing and organisational capability; its recognition of mapping as a lifecycle obligation; and the institutional fact that the three-domain thesis appears in an OECD framework.

What is not taken from it. The 35% claim as formulated; the proposition that mapping necessarily loses clinically relevant meaning; the standards table as a comparative capability assessment; the tripartite allocation as an optimal architecture; or the proposition that interoperability generally finances itself.

E-11: Interop Council (2026) — The Statutory Advisory Body Asks for a Duty on Semantic Data Holding

What it is. *Positionierung des Interop Council zu den Interoperabilitätsvorgaben im Rahmen des GeDIG*, two pages, Berlin, 11 August 2026, published on the national interoperability platform. The Interop Council is the body constituted under § 385 SGB V to advise the interoperability competence centre; seven full members are named, chaired by Sylvia Thun. The Council already appears in this assessment through its core-profile working group of March 2025 (D-10.3) and its working group on hospital-interface adoption of February 2024.

One disclosure, made here once rather than twice in places that do not know of each other. The chair of the Council is also deputy chair of HL7 Germany and the second author of the two Rinaldi and Thun publications appraised at D-8.3, whose first author is a member of her group at the Charité; the Charité belongs to the HiGHmed consortium whose use case those publications assess. **None of that touches the grade given here**, which rests on the wording of a published two-page position and on the positions it takes. It is recorded because the same person appears at two very different weights in this annex, once as the chair of a statutory advisory body and once as a co-author of a source at grade C, and a reader should not have to assemble that.

Evidence grade A for propositions about the wording of the paper, that being a proposition about a published document. **Evidence grade B** for the positions it takes, which are the published views of the body that holds them. No grade attaches to whether the positions will be enacted; the Council advises, the competence centre proposes, and a binding format is set by ordinance under § 385 SGB V.

Why it is carried here. The consultation statements of May 2026 divided on whether an interoperability obligation reaches the data that are exchanged or the model in which they are held (D-10.3). This paper is the first intervention in that dispute from the body constituted to advise on it, and it does not take either side: it divides the term the dispute was conducted with.

What it asks for, in its own words. A statutory duty of "standardisierte, semantisch-inhaltliche Datenhaltung auf logischer Ebene, sodass Daten standardisiert semantisch annotiert vorgehalten werden", paired in the same paragraph with the limit that "keine gesetzlichen Vorgaben für Datenbanken oder Systemarchitekturen gemacht werden" and that "die konkrete technische Implementierung bleibt in der Verantwortung der Systemhersteller".

The distinction it introduces. The debate recorded in D-10.3 ran between an obligation on exchange and an obligation on the storage model. The Council divides the second: the logical semantic layer is to be regulated, the technical architecture is not. That distinction is finer than either consultation position and is closer to the line this assessment draws than either pole, which is why it is set out rather than summarised.

Two limits the paper places on its own demand. The binding reaches not the logical layer at large but the part of it that has been determined: manufacturers are to have "keine semantischen Freiheiten hinsichtlich der Bedeutung oder Repräsentation von verbindlich festgelegten Dateninhalten besitzen, beispielsweise bei Maßeinheiten oder Beobachtungen". And the determined part is to begin small and grow, the European Patient Summary or a German specification derived from it serving as the basis for "eines ersten, verpflichtenden Datenkerns", with the paper recording that "für die praktische Umsetzung wird ein stufenweiser Ausbau befürwortet" and, because a patient summary does not carry every care situation, that "deshalb muss dieser Kern schrittweise erweitert werden". The demand is therefore that a named and initially small set of data contents be fixed in meaning, not that internal models be prescribed.

Why this bears on the reading of the consultation. It is tempting to record the Council as contradicting the vendor association, and that reading does not hold. The bvitg statement asks that internal data-holding models not be normed by statute and treats them as one thing; the Council divides that thing and disclaims half of it expressly. The two are incompatible only if bvitg meant the logical semantic layer as well, and compatible if it meant the store and the architecture. **The bvitg statement was read in its own wording on 18 August 2026** (56 pages; the sentence stands in the executive summary and again under § 386 SGB V), and it uses the term without defining it: nothing in

the document says whether "interne Datenhaltungsmodelle" means the store and the architecture, the semantic content, or both. The scope of the term therefore cannot be settled from the source, and neither reading is asserted here. What survives without a reading is that the dispute was conducted with an undivided term and that the advisory body is the first participant to divide it.

Three further propositions bear on this assessment.

- FHIR is to be "als semantischer und syntaktischer Standard das führende Austauschformat", with DICOM remaining admissible for imaging. This is the first occasion on which a § 385 body names the standard for the ordinance route rather than for a single programme (§ 1.2).
- Terminologies, codes and value sets are to be made binding across sectors and applications, LOINC and SNOMED CT being named as examples, and the Council concedes in the same passage that "eine vollständige Überprüfung im Rahmen von Konformitätsbewertungen nicht in allen Fällen möglich sein sollte". A body responsible for the assessment route conceding the limits of that route is unusual, and it supports from the other direction what D-10.8 records about the reach of conformity assessment.
- Data models themselves are to be standardised in statute in future, "aufgrund der Herausforderung, dass Terminologien von Datenmodellen abhängig sind", with DICOM and the aviation navigation standard ARINC 424 named as precedents. This is the sharpest sentence in the paper for the question of this assessment, and it is a proposal rather than a plan: no instrument, timetable or responsible body is named for it.

What the paper does not contain. openEHR is not mentioned, and neither is any other persistence standard. The body best placed to name one does not name one. That is recorded as an absence in a two-page document, not as a rejection, and no inference is drawn from it in either direction. The paper likewise names no version, profile family or artefact set to which the demanded semantic annotation would apply.

Interests. The Council is an advisory body at gematik, which is itself one of the specifying bodies whose fragmented mandates D-10.3 records. Its chair and members hold positions in German health informatics and standards work. The paper is a position and is presented as one.

E-12: Piera-Jiménez et al. (2026) — Spain's Ministry-Commissioned Delphi Names openEHR for the National Record

What it is. *Recommendations for a national electronic health record in Spain: a Delphi study*, Lancet Digital Health, 2026, DOI 10.1016/j.landig.2026.101037, open access. Commissioned by the Spanish Ministry of Health on 16 January 2023; a steering committee of 22 experts developed 45 statements in four domains, 220 experts nominated through the chief health information officers of all Spanish regions were surveyed, 152 responded (69.1%), and every item met the 70% consensus threshold in a single round, 44 of 45 at 80% or higher. Nine international advisors reviewed the resulting 20 recommendations; the external validation produced "no substantial changes". Funded by the EU (NextGenerationEU) through the Spanish Recovery Plan.

What it recommends, in its own words. The national record is to be built on the information architecture of ISO 18308:2011 (Rec 3.1). Compliance "is best achieved through dual information architecture models, specifically UNE-EN ISO 13606-1:2020 and openEHR models. Of these, openEHR is considered the most suitable for the development of the digital medical record of the NHS" (Rec 3.2). "For the exchange of information at the technical layer, the HL7 FHIR standard should be used, with its mapping defined in conjunction with the clinical knowledge model" (Rec 3.5). Further recommendations ask for progressive adoption that respects functioning source systems (Rec 3.3), a federated editorial governance for clinical information models led by the Ministry (Rec 4.1), standardised tooling (Rec 4.2), a central open repository of clinical information models (Rec 4.3), incorporation of the openEHR modelling methodology into Royal Decree 1093/2010 (Rec 4.4), and EHDS alignment (Rec 4.5).

Evidence grade A for propositions about the wording of the paper. **Evidence grade B** for the consensus figures, which are measured statements about what this panel agreed to. **No grade attaches to the suitability verdict itself**:

the paper states expressly that "the issued recommendations cannot be considered evidence-based", and a Delphi consensus is a structured record of expert opinion, not an outcome comparison. It is also a recommendation to the ministry and not an adopted national strategy; the status parallel is the Interop Council position at E-11, on the other side of the persistence question.

Why it is carried here. It is the weightiest institutional support the openEHR side has to date: national in scope, ministry-commissioned, internationally validated, published in a Lancet journal. Where E-2 to E-4 record positions of an industry vendor, a specialist society and an agency framework, this is a state-commissioned process for a country of 49 million. The set's symmetry rule requires that it be carried with the same care as the FHIR-side determinations, and it is the strongest named counter-position to this assessment's default for new infrastructure — institutional where Marfoglia et al. (E-9) is academic.

What it concedes without saying so. Rec 3.5 writes the mapping obligation of § 1.6 into the national plan: FHIR at the exchange layer with the mapping "defined in conjunction with the clinical knowledge model" is the Option D/E architecture of § 1.6.1 with its permanent, governed cross-standard boundary accepted in writing. The governance recommendations 4.1 to 4.3 — an editorial committee, shared tooling, an open model repository — are close counterparts of what Document 6 asks of whoever maintains a model holding.

Interests, and they are extensive. The first author is a volunteer board member of openEHR International and the lead author of the Catalonia chapter this assessment appraises at D-2; the declaration names a second author as CEO of openEHR International, and further authors hold openEHR roles (an unpaid openEHR member and co-developer of the Clinical Knowledge Manager prototype; consultants to openEHR International; a co-developer of the ISO 13606 series). HL7-side authors are present, among them an HL7 International and HL7 Italia board member, and the panel spans both camps; but the technically decisive authorship overlaps substantially with the leadership of the organisation whose standard the paper declares most suitable. The interests are openly declared over more than a column, the expert selection ran through the steering committee, and the paper names its own selection-bias risk. **A single-round unanimous consensus on 45 items is also a statement about how the questions were framed**, and the questionnaire was written by the steering committee.

One boundary of the introduction, checked against the primary source. The introduction states that "over 30% of information can be lost when exchanging data between systems from the same vendor, with this percentage escalating to 80% when different vendors are involved", citing Bernstam et al. 2022. The cited study (D-11.3) measured something narrower: the probability that a laboratory-test or medication *name string* arriving from one site had been seen at the receiving site, across 12 data elements at 68 US oncology sites — 0.68 within a vendor, 0.22 between vendors. Its authors call the element set "an example, rather than a definitive list" and the score a first attempt, and the study says nothing about information being lost in exchange, about FHIR or about openEHR. The 30/80 sentence reads a recognition score as a loss rate. This bounds the introduction's framing, not the recommendations.

What is taken from it. The institutional fact and its wording; the consensus figures; the written acceptance of the mapping obligation; the governance recommendations as convergent with Document 6.

What is not taken from it. The suitability verdict as evidence of comparative fitness; any implication that the recommendation is enacted; the 30/80 loss claim.

Appendix G: Source Ledger

This ledger records, for every load-bearing source, the basis on which it was assessed, the proposition it supports, its grade, relevant interests and retrieval status.

Inclusion rule, and the reason for it. A source is listed here when a specific finding or graded statement relies on a **judgement** about it: how far it can be trusted, whose interest it serves, whether it establishes what it is cited for. That is what this ledger is a register of.

Statutes, normative specifications, implementation guides, registry manifests and recorded bibliometric queries are registered at Appendix H; where a graded proposition rests on a reading of one of them, the row appears here as well and Appendix H says so. Such a reading contains no judgement to record: the artefact either says the thing or it does not, and a reader checks it by opening it.

Those artefacts are not omitted from the audit trail. They are listed in Appendix H, the primary artefact register, with their location and retrieval date and without an evidence grade. Unpublished working materials and background sources appear in neither register where no finding depends on them.

Basis of assessment.

- **Full text:** the publication was read in full.
- **Full text + extract:** the publication was read in full and a model-produced summary is also on file.
- **Extract:** the summary extract was the working basis and the publication was not read in full.
- **Artefact inspected:** the cited specification, guide, position paper, documentation or circular was inspected directly.
- **Artefact inspected + extract:** as above, with a model-produced summary also on file.
- **Counted:** the figure was produced by the author from identified artefacts, with the query or count recorded in Appendix H.
- **Reported only:** no accessible primary text was available.
- **Personal communication:** the evidence rests on unpublished communication to the author.

The summary extract accompanies this document as Literature Summaries.md. **Two limits on the "Full text" label, and they are stated rather than smoothed over.** It indicates that the complete publication was available and was read. It does **not** indicate that the study was independently replicated, that its underlying data were checked, or that its analysis was reproduced, only that the text, and not a summary of it, was the working basis.

And a number of rows carry a basis of assessment that was **reconstructed analytically from how the source is used in the text, rather than reported by the author**. Those rows are the ones a reader should weigh most carefully, and correcting them against the author's own recollection is an outstanding task recorded in the project handover rather than a completed one. The other labels identify an artefact or a record that can be checked separately, which is why they do not carry the same reservation.

Grading rule. A grade attaches to a proposition-source pair, not to a publisher or an entire document. The assessment proceeds in four steps:

1. Classify the proposition as a requirement, feasibility claim, usage claim, effectiveness or cost claim, or institutional position.
2. Determine whether the source constitutes the proposition or merely reports it.
3. Determine whether the assertion can be checked at the cited location.
4. Record whose interest the proposition serves and lower the grade where an interested source is the sole support.

Grades express how strongly the named source supports the bounded proposition stated in the ledger. They do not convert feasibility into effectiveness, association into causation or absence from this search into proof of non-

existence. Claims that no study, product or implementation was found are therefore stated as search results and are not assigned the grade of the sources that happened to be located.

GRADE	MEANING
A	Constitutive or directly repeatable evidence: statutes, normative requirements, register entries and reproducible counts from identified material.
B	Independently inspectable but not constitutive evidence, including official frameworks and published working-group documents read for what they state.
C	Reported evidence without a controlling interest, including press-reported usage figures and peer-reviewed viewpoints or short reports.
D	Evidence reported by an interested party or not independently evaluated, including vendor documentation, implementation presentations, blogs, posters, preprints and personal communication.

A normative specification receives grade A for what it prescribes, not for implementation in practice. A statement against the apparent interest of its source is recorded as context and does not raise the grade.

Retrieval terms, and what they do not say. *Open* means freely accessible. *Paywalled* requires payment or institutional access. *Restricted* is available only to a limited audience. *Unpublished* is held by the author. *Versioned* identifies the relied-upon state. *Volatile* may change without an archived copy. *Live aggregate* changes with reporting sites and query date. **These terms describe access, not permanence.** A source can be open today and moved tomorrow, and "open" says nothing about whether the same state of the same document can be reached again.

Permanence is carried by the **persistent identifier in the first column** — a DOI where the publication has one, otherwise a canonical or repository address. Where a row has no identifier, that is itself the finding: the source was read but cannot be pointed at, and the grade already reflects it. Web artefacts without a persistent identifier carry their retrieval date in the row that cites them or in the appendix that discusses them.

Grades marked † are assigned in this ledger where no grade was previously stated. The ledger records the assessment process but does not increase the evidential value of a source.

G-1: Peer-Reviewed Literature and External Analyses

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Piera-Jiménez et al. 2026, Delphi study, Lancet Digit Health 10.1016/j.landig.2026.101037	Full text	Spain's ministry-commissioned recommendation of openEHR for the national record with FHIR at the exchange layer, the consensus figures, and the written acceptance of the mapping	A for the wording; B for the consensus figures; no grade for the suitability verdict	Authorship overlaps with openEHR International leadership, declared in the article over more than a column; panel selected by the steering committee; EU-funded	Open access

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		obligation (E-12, § 1.1, § 1.4, § 2.4.1)			
Bernstam et al. 2022, JAMIA 10.1093/jamia/ocab289	Full text	The measured interoperability scores 0.68 intra-vendor and 0.22 inter-vendor over 12 structured elements at 68 US oncology sites, and their boundaries (D-11.3, § 1.5)	B for the scores; no grade for FHIR, openEHR or loss claims	Three authors contributed as CancerLinQ LLC employees, disclosed	Paywalled; full text held by the author
Pedrera-Jiménez et al. 2023, JMIR 10.2196/48702	Full text + extract	Compatibility of openEHR, ISO 13606 and FHIR under the agnostic approach; the dual-qualification staffing requirement it does not conceal (§ 1.4, § 1.6); the Madrid INFOBANCO case	C †	Authors include openEHR-side figures (Kalra, Beale) and the implementing hospital	Open
Christensen & Ellingsen 2016, Int J Med Inf 10.1016/j.ijmedinf.2016.02.004	Full text + extract	A longitudinal case study of the FIKS programme and its large-scale model-driven implementation challenges (D-3, D-5)	A, with the stated qualification that it examines one project rather than comparative architectural fitness	None declared; independent of both camps	Paywalled
D'Amore et al. 2014, JAMIA 10.1136/amiajnl-2014-002883	Full text	Semantic errors and permitted variation in 21 manually inspected C-CDA documents, selected from 91 documents produced by 21 technologies;	B † for the inspected artefacts and reported observations; no grade for FHIR, architecture	ONC-funded; authors drawn partly from consultancies supplying document-quality tooling	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		optionality was the largest category (§ 1.6.2; D-11.1)	or patient harm		
Wright et al. 2018, J Gen Intern Med 10.1007/s11606-018-4415-9	Full text	That alert frequency and acceptance changed substantially after one system replacement, illustrating that data portability does not ensure functional equivalence (§ 1.6.2; D-11.2)	A † for the measurements at that site; no grade for causation, patient harm or architectural comparison	NLM and AHRQ funded, no conflicts declared; one institution and one product installation	Open
Bruthans & Jiráková 2023, J Med Syst 10.1007/s10916-023-01920-9	Full text + extract	A published KPI analysis of eHDSI and its record of the programme's largely test character during the observation period (D-4)	C †	None declared	Open
Heyder et al. 2023, Bundesgesundheitsblatt 10.1007/s00103-022-03649-1	Full text + extract	The NUM research-data-platform architecture, and by what it does not describe, the absence of the openEHR side (C-3, D-4)	B †	Written by the platform's own consortium	Open
Sass et al. 2020, BMC Med Inform Decis Mak 10.1186/s12911-020-01374-w	Full text	The paper's own record that the MII FHIR profiles were taken into account in defining GECCO, which dates the direction of influence (C-3)	B †	The present author is a co-author, disclosed at the point of use	Open
Ladas et al. 2022, Stud	Full text	That COVID data held in openEHR	B	Authors from the Hannover	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Health Technol Inform 10.3233/SHTI210963		repositories had to be converted into GECCO FHIR for exchange (C-3)		openEHR environment	
Behrend & Ingenerf 2022, GMDS meeting abstract, not a full paper 10.3205/22gmds069	Full text of the abstract	That GECCO profilings exist in both FHIR R4 and openEHR and both were held equally applicable (C-3)	C — a conference abstract records that the work was presented, not that it was reviewed as a comparative study	None declared	Open
Hyppönen et al. 2019, Int J Med Inf Julkari repository	Extract; bibliographic record verified directly	The Finnish usage row of Appendix D-4, resting on national physician surveys of 2010, 2014 and 2017	C	Authors from Finnish national bodies	Paywalled
Piera-Jiménez & Carot-Sans 2025, book chapter, ch. 8 in a Springer volume 10.1007/978-3-031-80704-6_8	Full text + extract	The Catalonia procurement record (§ 1.4.1, D-2, D-3)	C †	The first author is now CEO of openEHR International; the chapter is an institutional self-description	Paywalled
Borys et al. 2025, JMIR Cancer 10.2196/73691	Full text + extract	The Essen dermatology dashboard as a FHIR production case (D-7.3)	C	Authored by the implementing group	Open
Schmidt CS et al., Transfus Apher Sci 65 — published	Full text + extract	The Essen blood-product dashboard as a FHIR production case (D-7.3)	C	Authored by the same implementing group	Paywalled

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
online 2025, issue dated 2026 10.1016/j.transci.2025.104301					
Ammon et al. 2024, Bundesgesundheitsblatt 10.1007/s00103-024-03888-4	Full text + extract	The MII core dataset and the state of its FHIR-based integration infrastructure (D-7.4)	B	Authored by the initiative's own interoperability working group	Open
Kohler et al. 2025, arXiv preprint arXiv:2511.14618	Full text + extract	The FHIRConnect mapping specification and the openFHIR engine as the concrete tooling of the mapping layer (§ 1.6)	D	Authors include openEHR International figures; the reference engine is offered commercially as well as open source	Open, not peer-reviewed
Kohler et al. 2023, J Biomed Inform 10.1016/j.jbi.2023.104437	Full text	The declarative archetype-to-OMOP mapping language OMOCL and the Eos tool that executes it, and the authors' own conclusion that transforming openEHR data into the less expressive OMOP CDM loses semantics (§ 2.7)	B †	Written by the group that built the bridge; the loss statement runs against the interest of that work and is weighted accordingly	Open access
Lenert et al. 2021, JAMIA 10.1093/jamia/ocab108	Full text	That a repository holding clinical data as FHIR canonically can produce PCORnet and OMOP research marts automatically, with FHIR Subscription	C †	Reported by the implementing group; two authors are employed by a FHIR server vendor (Smile CDR) — disclosed at the point of use	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		keeping the delay small (§ 2.7)			
Essaid et al. 2024, JAMIA Open 10.1093/jamiaopen/ooae045	Full text	That an existing OMOP warehouse can be transformed into US-Core-conformant FHIR resources and exported through Bulk FHIR, at a non-conformance rate below one percent (§ 2.7)	C †	Reported by the implementing consortium. The resource count given in the abstract is implausible by two orders of magnitude and is deliberately not relied on here	Open access
Jayathissa et al. 2025, Stud Health Technol Inform 10.3233/SHTI250432	Artefact inspected	That an ETL from FHIR Bundles into the OMOP CDM is described as a worked route (§ 2.7)	D †	Short conference contribution by the implementing group, no independent evaluation	Restricted, conference proceedings
Kapitan et al. 2025, J Med Internet Res 10.2196/66616	Full text	The peer-reviewed counter-position to the three-domain thesis (E-1)	C †	Authors argue from the open-source implementation environment	Open
Delussu et al. 2024, survey of openEHR CDRs 10.1016/j.ijmedinf.2024.105591	Full text	The openEHR repository product layer as declared by its vendors: seventeen production-capable repositories identified and eleven surveyed; AQL in 6 of 11 and SQL or a vendor formalism in the other 5; a mostly conformant openEHR REST API in 4; the openEHR EHR Extract in 1,	B for the declared product properties, which are reproducible from the published answer set; C for any inference from a declared property to what a deployment can do; no grade for	Vendor self-declaration throughout, unverified by the authors; 35% of suppliers approached did not respond; the author team builds openEHR tooling, which strengthens the findings running against openEHR's claims and weakens the favourable ones	Public, peer-reviewed; answer set and analysis code published by the authors; data current to 2023

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		plus 1 through a separate product; free and open source in 2; ADL 2.0 in 1; and dispersed Composition serialisations, which the authors themselves say prevent straightforward interoperability between the repositories in the absence of a conformance testing framework (§ 1.5, § 2.6.1, § 2.6.2, A-2)	any comparison with the FHIR product layer, whose source here is weaker		
Kramer 2023, AMIA Annu Symp Proc PMC10148270	Full text	That FHIR's intended 80/20 rule is closer to 65/35 across 2,897 profiles from 125 implementation guides: base FHIR supplied approximately two thirds of the elements the guides studied needed, profiling defines 33 to 38% of new data elements, two profiled items (Extension and Observation) account for 49% of all profiles, and of 2,811 standalone extensions 2,283 were custom. Used in E-10 to show what the OECD's paragraph	B for the proportions reported over the identified sample, which is a primary study of published artefacts; no grade for inference to all FHIR implementations, and none for the claim that profiled elements cannot be shared, which the study does not make	MITRE, the same institution as D-14; the finding is unfavourable to the standard this assessment recommends	Public, peer-reviewed conference proceedings, open access through PubMed Central

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		20 cites and what it does not establish			
Kramer & Moesel 2023, JAMIA Open 10.1093/jamiaopen/ooad001	Full text	That conformance of two systems to their respective specifications does not imply they can interoperate, computed on US Core 5.0.1 against IPS 1.0.0: Patient and Condition potentially infeasible, MedicationRequest infeasible; and that only a `required` binding definitively limits the set of instances a profile admits, extensible being treated by tooling as unconstraining (§ 1.5, § 2.5.2, D-14; Document 5)	A for the relationships computed between the two named profile versions, which are direct readings of published artefacts; B for the method as a general instrument, proposed and demonstrated but not independently evaluated; no grade for prevalence, three of twenty-three resource types being reported and not sampled	MITRE, which builds FHIR tooling for US federal programmes and maintains two of the validators the paper names; no competing interest declared, no external funding. The finding runs against the standard this assessment recommends	Public, peer-reviewed
Jones et al. 2021, JAMIA 10.1093/jamia/ocab028	Full text	That of the organisations building the SMART/HL7 bulk data interface ahead of the compliance date, only the electronic-health-record vendors named regulation	C — a self-selected survey of 22 respondents from 37 approached, self-reported intentions rather than measured deployments	Authors include the SMART team that wrote the specification surveyed and an HL7 officer; the subject is uptake of their own work	Public, peer-reviewed

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		as the primary motivation, the payors, cloud vendors, research institutions and tooling developers citing use cases of their own (§ 1.2, point 2)	, United States, 2020		
Mandl et al. 2020, npj Digital Medicine 10.1038/s41746-020-00358-4	Full text	That custom, site-specific mapping into i2b2, PCORnet and OMOP confines participation to large academic medical centres with specialised and often grant-funded teams — the mapping-cost proposition of § 1.6 stated by an interested party on the other side of the debate	C — an architecture and advocacy paper in a peer-reviewed journal; the proposition is asserted from the authors' experience, not measured	Authors include the FHIR product director and HL7; the paper advocates the API it describes	Public, open access
Griffin et al. 2022, JAMIA Open 10.1093/jamiaopen/ooac077	Full text	The application-layer inventory of 112 real-world FHIR applications, of which 49 appear in no electronic-health-record application gallery (A-1)	C — a brief communication with a convenience sample and self-reported data	Several authors are officers of HL7	Public, open access CC BY

G-2: Systematic Reviews and the Sources Drawn From Them

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Vorisek et al. 2022, JMIR Med Inform 10.2196/35724	Full text + extract	The pre-registered review of FHIR use in health research (D-8.3); the terminology-	B	The last author is deputy chair of HL7 Germany, declared in the paper and disclosed in D-8.3	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		binding rates used in § 2.5.1; the modelling-depth finding of § 2.4			
Tabari et al. 2024, JMIR Med Inform 10.2196/58445	Full text + extract	The scoping review (D-8.2), including the share of dynamic pipeline models among the implementations examined	B/C	None declared	Open
Ayaz et al. 2021, JMIR Med Inform 10.2196/21929	Full text + extract	The inventory of FHIR implementations and challenges (D-8.4), usable as an inventory rather than as an argument	C	None declared	Open
Rinaldi & Thun 2021, two publications: Stud Health Technol Inform 10.3233/SHTI210189 **and GMS Med Inform Biom Epidemiol 17(2):Doc07**	Full text of both	The documented openEHR to FHIR to OMOP transformation path (§ 1.6, § 2.4.1); from the journal article the ISO/TR 12300 equivalence scores and the coverage figures (§ 2.7, § 2.6)	C †	The second author is deputy chair of HL7 Germany **and chairs the Interop Council appraised at E-11**; the first author is a member of her group at the Charité, which belongs to the HiGHmed consortium whose use case is assessed here. The quality assessment is therefore a self-assessment by the implementing party, as with Hund et al. and Wettstein et al. below	Open
Hund et al. 2021, Stud Health	Full text + extract	The HiGHmed Data Sharing Framework	C †	Authored by the implementing consortium	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Technol Inform 10.3233/SHTI210060		exchanges process information using FHIR R4 resources; it does not alone establish an openEHR→FHIR clinical-payload path (§ 1.6.1)			
Wettstein et al. 2021, *Feasibility Queries in Distributed Architectures*, Stud Health Technol Inform 278:134–141 10.3233/SHTI210061	Full text + extract	A FHIR-based distributed feasibility-query implementation within HiGHmed, not a complete classification of each site's persistence path (§ 1.6.1)	C †	Authored by the implementing consortium	Open at the DOI; the publisher page is not reliably free
Wettstein et al. 2021, *Data Sharing in Distributed Architectures*, Stud Health Technol Inform 283:111–118 10.3233/SHTI210548	Full text + extract	The FHIR-based process data model and sharing layer of the HiGHmed framework (§ 1.6.1)	C †	Authored by the implementing consortium	Open at the DOI; the publisher page is not reliably free
Finster M, Wenzel M, Taghizadeh E (2025), *Common data models and data standards for tabular health data: a systematic review*, BMC Med Inform	Full text	The systematic comparison of four common data models against four data standards for tabular health data, in which OMOP and FHIR score highest overall, together with the conclusion	B †	No interest declared. The criteria include popularity and support, so a high score is in part a measure of adoption; the scope is tabular health data	Open access

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Decis Mak 25:422, 13 November 2025 10.1186/s12911-025-03267-2		that no single global representation can be selected (§ 2.7)			
Duda et al. 2022, JAMIA 10.1093/jamia/ocac105	Full text	The scoping review of 203 FHIR projects for clinical research and its distribution, in which mapping data into and out of FHIR is the single largest category at 32.5% (D-8.5, D-8.6)	B/C — a documented systematic search counting self-described project purposes that were not independently verified	Two authors hold roles at HL7 Da Vinci and the National Library of Medicine; the work was funded under an NLM contract to promote FHIR for research, declared in the paper	Public, peer-reviewed
Marfoglia et al. 2026, systematic review 10.1016/j.jbi.2026.105022	Full text	The distribution and adaptation of FHIR, OMOP-CDM and openEHR across 99 peer-reviewed studies from 2021 to 2024, the 85% adaptation rate for FHIR, the 27% coverage-reporting rate, and the hybrid-ecosystem conclusion carried as a counter-position (§ 1.2, § 1.5, § 2.7, E-9)	B for the counts over the review's own inclusion list, which are propositions about a body of literature; no grade for any inference from that literature to deployment, which the review itself declines	Academic team, no declared standards affiliation, national research funding; large language model disclosed for proofreading	Public, open access CC BY

G-3: Published Professional-Society and Camp Materials

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Oemig, Blobel, Helmer, Birkle 2018, HL7-Mitteilungen 41 issue PDF	Full text + extract	The distinction between conformance and interoperability as orthogonal properties (§ 1.5)	C †	Published in the periodical of HL7 Germany, editorially edited and not peer-reviewed; flagged as FHIR-leaning in "On methodology and sources"	Open
Blobel & Oemig 2018, HL7-Mitteilungen 41 issue PDF	Full text + extract	The seven-level interoperability grid extending the HIMSS classification (§ 1.1, § 2.1)	C †	Same periodical and same status	Open
Pazos et al. 2015, tutorial presentation openEHR tutorial page	Full text + extract	The persistence statement of the openEHR source itself, and the slide observations of Appendix B (§ 2.2)	C †	An internal openEHR source by origin, not by publication status	Open
Beale, 19 August 2020, openEHR Discourse, thread *Addition of required / extensible / preferred / example terminology constraint to AOM and ADL*	Artefact inspected	The proposal to add FHIR's four binding strengths to AOM and ADL, with FHIR named as the template (§ 2.5.2)	D, weighted higher as a statement against the interest of the camp making it	Specifications Editorial Committee of openEHR International	Open, dated; no visible conclusion to the thread

G-4: Analyst, Position, Framework and Institutional Documents

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Gartner Research Note, April 2017	Reported only	The external industry-analytical counterpoint of	C/D †	Independent of both camps, but a commercial analyst	Restricted; no accessible full text was obtained. Existence, title and

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		Appendix D-6: the recommendations, the three suitable scenarios and the cautions, in each case as summarised rather than as published		product without peer review	date are independently corroborated by the European Commission staff working document SWD(2018) 126 final, footnote 51; the author names and the document number are not
Tsafnat et al. 2024, J Med Internet Res 10.2196/55779	Extract	The three-domain thesis in its most prominent formulation (E-1)	D	Authors are institutional representatives of the standards recommended; conflicts declared in the article	Open, editorial, not peer-reviewed
Schabetsberger & Mangesius 2026, white paper	Artefact inspected	The comparative vendor analysis of E-2, including the modelling-phase duration reused by the ÖGTelemed paper	C/D	Vendor paper, self-declared not peer-reviewed	Open
ÖGTelemed 2026, position paper v1.0	Artefact inspected	The counter-finding that both standards can serve vendor-independent persistence (E-3, § 1.2)	C	Position paper of a specialist society, not of a chamber or a public authority	Versioned, 24 June 2026
Australian Digital Health Agency 2026, national framework agency PDF	Artefact inspected	Governance statements on standards selection at national level (E-4)	B	Official framework of the responsible agency	Open
Fuhrmann et al. 2026, conference poster	Artefact inspected	The CKM consensus-review process as a governance illustration (E-5)	D	Produced within the openEHR eye-care collaboration	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Devitt, 28 May 2026, blog post	Artefact inspected	The "Red Hat problem" as a distinct lock-in mechanism (§ 1.3 Dimension 1, E-6)	D	Single openEHR-leaning author, no peer review, no source references. The statement runs against the apparent interest of its author, which lowers the likelihood of simple partisanship but is not itself evidential weight and does not raise the grade; the premise of technical superiority is expressly not adopted	Volatile
Devitt 2026, FHIR server catalogue	Artefact inspected	Market breadth on the order of thirty FHIR server products (§ 1.3 Dimension 1, Appendix A-1)	D	Same author and same status	Volatile, updated July 2026
openEHR Switzerland 2026, guest post	Artefact inspected	The Swiss joint working group and the vaccination-showcase blueprint (C-1, E-8)	D	Camp self-presentation of a hackathon proof of concept	Open
GeDIG cabinet draft 2026	Full text	That § 386a SGB V obliges manufacturers to release health data in an interoperable format and to enable access, the format set by ordinance; that § 386c prohibits formats preventing complete	A	Government draft; no interest of its own in the architecture question	Published and open; supplied to the author as a file on 4 August 2026 after four retrieval attempts had failed from the working environment

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		semantic reconstruction and the withholding of elements needed for documentation and archiving; and that the provider-side holding duty appears only in the summary sheet and the general memorandum, with no operative provision and no amendment to § 386 (D-10.3, § 1.6.3)			
OECD 2026, *Interoperability in Healthcare: Towards an Interconnected Future* 10.1787/eb71b7be-en	Full text	The multi-dimensional policy framing of interoperability; reported governance, standards-adoption and conformance barriers; mapping as a lifecycle obligation; and the appearance of the three-domain thesis in an international policy framework (§ 1.5, § 1.6, § 2.4.1, E-10)	B for the OECD framework and its reported survey findings; C for third-party quantitative projections; no grade for the technical correctness of its illustrative standards taxonomy	Intergovernmental policy analysis based on interviews, 22 country responses, consultation and desk research; no declared standard or product interest. Several standard-specific descriptions depend on standards organisations or an industry blog	Public, open; OECD Health Working Paper No. 197, published 24 July 2026
Health Samurai 2026, FHIR server benchmark	Artefact inspected	Throughput, latency and storage figures for Aidbox, HAPI FHIR, Medplum and the	D — the publisher builds the product that leads on every	Vendor benchmarking its own product against competitors	Public; code, harness and results published on GitHub, report regenerated daily;

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
report · harness		Microsoft FHIR Server on one machine against 1,000 Synthea patients and roughly two million R4 resources (A-1)	axis but storage; the dataset is far below deployment scale despite the title; a single configuratio n was measured. The published harness is why it is listed at all and does not raise the grade		retrieved 13 August 2026
CTSA data-harmonizati on, FHIR repository evaluation project page	Artefact inspected	That seven FHIR repository products were compared across seven criteria by a research network, with no overall recommendation reached (A-1)	D, and stale — the evaluators state they set each server up from the available documentat ion without expert consultatio n, and the reference points are HAPI 3.8.0 and 4.0.0, placing it in 2019	A CTSA project comparing third-party products; no product interest apparent	Public project documentation; retrieved 13 August 2026
Joint Initiative Council, member list jointinitiativ ecouncil.org	Artefact inspected	That the JIC has eight member organisations — CDISC, CEN/TC 251, DICOM, GS1, HL7, IHE, ISO/TC 215 and SNOMED	B — the published description of a partner organisatio n, consistent	Reading of the published member list	Public

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		— and that it is a liaison and coordination body rather than an accreditation body (§ 1.2, point 5)	across the JIC's own site and CDISC's; retrieved 12 August 2026		
bvitg, DKG and ZVEI consultation statements, May 2026	Artefact inspected	The positions of the vendor, hospital and device associations on §§ 386 and 386a, and the contested line between interoperable exchange and internal data holding (D-10.3)	B	Association statements, each representing the interests of its members	Open, dated
Interop Council, *Positionierung zu den Interoperabilitätsvorgaben im Rahmen des GeDIG*, 11 August 2026	Full text	The statutory advisory body's demand for semantically standardised data holding at the logical layer, bounded to bindingly determined data contents, and for FHIR as the leading exchange format (E-11, § 1.2, § 1.6.3, Document 5 D-10.3)	A for the wording; B for the positions	Advisory body at gematik; chair and members active in German standards work	Open, dated 11 August 2026, retrieved 18 August 2026

G-5: Normative Specifications and Published Implementation Guides

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
openEHR International, Data Types Information Model	Artefact inspected + extract	That `terminology_id = "local"` means the origin of the permissible	A	Published by openEHR International; the proposition concerns what the	Versioned, latest release

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		values is described within the archetype (§ 2.5.1)		specification requires	
openEHR International, Architecture Overview ch. 12, BASE Release 1.0.2	Artefact inspected + extract	The double function of local at-codes, the definition of the archetype path, and the specification's own caveat that external terminology coverage is usually incomplete	A	Same publisher, same character of proposition	Versioned, latest release
openEHR International, Support Terminology Specification	Artefact inspected + extract	The delimitation of the openEHR-internal support terminology from external clinical terminologies (§ 2.5.1)	A	Same publisher, same character of proposition	Versioned, Release 3.0.0
freshEHR / Digital Health Innovation Scotland, implementation documentation	Artefact inspected + extract	That in the simplified FLAT format value and terminology may be omitted because the template fixes them, and that the operational template is the authoritative validation source	C	Vendor-affiliated documentation from the openEHR environment; the statement runs against the apparent interest of its source, which is recorded as an interest observation and not as a grade uplift	Volatile; no archived copy is held, and the pages can be revised
HL7 International, FHIR Data Types `Coding` and `CodeableConcept`, R4 and R5	Artefact inspected + extract	That `Coding.system` is `0..1` and that a `CodeableConcept` carrying `text` alone is valid, hence the delimitation in § 2.5.1	A	Published by HL7 International; the proposition concerns what the specification requires	Versioned, R4 and R5

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
HL7 International / Clinical Interoperability Council, FHIR R4 Symptoms IG, STU 1	Artefact inspected + extract	That the FHIR Logical Model holds authoritative status while the openEHR archetype is carried as an informative document (E-7)	B	Working-group guide of one of the two standards bodies under comparison	Versioned, build April 2026
HL7 International, FHIR Structured Data Capture Implementation Guide	Artefact inspected	That SDC extends `Questionnaire` and `QuestionnaireResponse` with rendering, pre-population, calculated expressions, conditional display, modular forms and extraction into clinical resources by way of `StructureMap` (§ 2.6.1 (derivation at D-15))	A for what the guide specifies; no grade for its uptake	Working-group guide of one of the two standards bodies under comparison	Versioned, published guide
HL7 International, FHIR QuestionnaireResponse search parameters, R4, R5 and current build	Artefact inspected	That R4 defines no search parameter on item answers or `linkId`, that R5 adds only `item-subject`, and that the current build adds composite parameters pairing `linkId` with the answer value (§ 2.6.1 (derivation at D-15))	A for what each release defines; no grade for server support in the field	Published by HL7 International; the proposition concerns what the specification defines	Versioned, R4, R5 and continuous build

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
HL7 International, SDC Modular Forms and the `$assemble` operation	Artefact inspected	That SDC specifies composition of Questionnaires — a `subQuestionnaire` extension naming another Questionnaire by canonical URL, recommended version-specific, expanded recursively by `$assemble` — while establishing no curated library of form modules (§ 2.6.1 (derivation at D-15))	A for what the guide specifies; no grade for the existence of any library	Working-group guide of one of the two standards bodies under comparison	Versioned, published guide
AKTIN, emergency-department FHIR Implementation Guide	Artefact inspected	That register reporting is organised as Questionnaires pre-populated from a care-based record assembled from ordinary clinical resources, with derived items computed rather than separately documented and items annotated with LOINC, SNOMED CT and ATC (§ 2.6.1 (derivation at D-15); D-13)	B †	Published through the ART-DECOR environment, in which the author is involved — disclosed at the point of use in D-13	Open, versioned publication
REDCap Shared Library, Vanderbilt University Medical	Artefact inspected	That a curated library of data-collection instruments exists outside either standard, reviewed by an	C †	Operated by the organisation that develops REDCap; access is restricted to authenticated users at partner institutions, so the	Restricted

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Center library		oversight committee for relevance, correct functioning and copyright before release, and importable into projects (§ 2.6.1 (derivation at D-15))		holding could not be inspected item by item	
Xt-EHR / European Commission, EHDS Logical Information Models	Artefact inspected + extract	That logical models are rendered by the ordinary FHIR publication toolchain, and that the UML diagrams are authored alongside the models rather than generated from them (E-7)	B	Guide of an EU-funded joint action	Versioned
ART-DECOR Open Tools, product documentation	Artefact inspected + extract	The export asymmetry by which a dataset carries no cardinality and exports with `0..*` while a transaction transfers its cardinality into the logical model (E-7)	D	The present author is involved in the ART-DECOR Expert Group; the interest is disclosed in E-7	Volatile, product documentation
openEHR International / HL7 International, Joint Statement May 2026	Artefact inspected + extract	The record of the second joint annual meeting, read as institutional self-presentation (E-8)	D	Joint self-presentation of both organisations	Open
HL7 Europe / IHE-Europe, EURIDICE	Artefact inspected	That the specifications intended for	B †	The author is CEO of HL7 Deutschland and Senior Expert at	Open, volatile

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
specification catalogue and coordinated ballot announcement		reference in the EHDS Implementing Acts are being produced as FHIR implementation guides by HL7 Europe and IHE-Europe, and that two of them were in coordinated ballot to 30 April 2026 (§ 1.1, D-12)		HL7 Europe, one of the two publishing bodies — disclosed at the point of use in D-12. The proposition is a reading of published artefacts and checkable against them	
HL7 FHIR specification, binding strength	Artefact inspected	That `ElementDefinition.binding.strength` distinguishes required, extensible, preferred and example, and that only required compels membership in the value set (§ 2.5.2)	A	Normative specification of the standards body	Open
openEHR AM/ADL2 specification, term_bindings	Artefact inspected	That ADL and the AOM define an optional `term_bindings` sub-section mapping at-codes to external terminology codes, without a strength (§ 2.5.2)	A	Normative specification of the standards body	Open

G-6: Reported Production Cases and Programme Sources

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Welsh Government 2023, Welsh Health	Artefact inspected	The binding obligation on all NHS Wales bodies to adopt	A for what the circular requires	Government circular of the responsible administration	Open

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Circular WHC/2023/018 gov.wales		HL7 FHIR as a foundational standard (D-7.2)			
Digital Health and Care Wales, Wales FHIR Implementation Guide Simplifier	Artefact inspected	That the Welsh profiles are maintained as a national guide derived from UK Core (D-7.2)	B	Published by the implementing body	Versioned, v2.6.0; the publishing host was not reachable from the working environment, and the guide is cited by its published identifier and version
HL7 UK, 4 June 2026, trade article *Welsh FHIR: Progressing interoperability across Wales*	Artefact inspected	The Welsh Care Data Repository as a persistent native FHIR store (D-7.2)	C †	Trade article with a DHCW representative; institutional self-presentation without independent evaluation	Open
Namlı 2025, DevDays presentation	Reported only	The Turkish national repository reported at some 38 billion FHIR resources on around 32 million patients (D-7.5)	D	Presentation by the implementing party, no independent evaluation	Restricted, conference slides and recording
Manzelli 2025, DevDays presentation	Reported only	The Lombardy regional clinical data repository at around 10 million patients (D-7.5)	D	Presentation by the implementing party	Restricted, conference slides and recording
Chai 2024, DevDays presentation	Reported only	The Singapore national programme case, which does not determine the persistence model (D-7.5)	D	Presentation by the implementing party	Restricted, conference slides and recording

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Smithies 2026, personal communication	Personal communication	The Ukrainian figures and the qualification, stated by the reporting source itself, that the REST interface there is not a FHIR API (D-7.5)	D	Communication from a source active on the HL7 side; no published primary source	Unpublished
Medical Informatics Initiative / TMF, FDPG holdings counter	Artefact inspected	The aggregate data holdings of the participating data integration centres (D-7.4, Finding 3)	C	Self-reported by the participating centres, publicly auditable but not independently verified	Live aggregate, figures as of 22 July 2026
Medical Informatics Initiative, core-dataset implementation guide *Manteldokument* (Simplifier.net) and the core-dataset page on medizininformatik-initiative.de	Artefact inspected	The dating of the two decisions behind the initiative's FHIR adoption: the textual core dataset of March 2017 and the National Steering Committee decision of April 2019 for HL7 FHIR as the technical representation, with ART-DECOR, Forge and Simplifier.net as tooling (C-3)	B	The initiative's own account of its own decisions	Open, retrieved 18 August 2026
gematik, Competence Centre for Interoperability — conformity assessment (KOB)	Artefact inspected	That a mandatory conformity-assessment procedure with supplied test suites, issued confirmations and a public positive list	A for the statutory obligation; B for the published description of the procedure; D † for the figures,	The body that runs the procedure reporting on its own procedure. ISiK, one of the announced subjects, is developed in proximity to the German HL7	Open, volatile; figures as of 11 December 2025

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		operates under §§ 385 and 387 SGB V, and its reported state of 141 assessed systems at some 98% of the practice-software market (Document 5, D-10.8; § 1.7)	which are self-reported by the operating body	community, which the author leads — disclosed at the point of use	

German market sources (Appendix D-7.7 to D-7.9):

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
Tiplu GmbH, TipluDB product page	Artefact inspected	That the repository is fed from the HIS and subsystems, and its data scope — the basis for classifying the product as Option B (D-7.7)	D †	The vendor of the product described	Open, volatile
Tiplu, interview in the HL7 Germany magazine	Artefact inspected	That FHIR is used for analysis and long-term storage rather than communication alone, and that the FHIR server was built in-house for search performance (D-7.7)	D †	Company representative on the company's own architecture, in a publication of the standards organisation the author leads — disclosed at the point of use	Open, volatile
e-health-com, company news on the funded project	Artefact inspected	The funding, the term, the EU-HIP context, and the stated aim of collecting treatment data as FHIR data without additional	C † for the funding and its framing; D † for the architectural aim, which traces to	Company-supplied press item in a trade outlet	Open, volatile

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
		transformation (D-7.7)	the company		
medconweb, report on the funded project	Artefact inspected	The funding body, the project term and the managing director's statement on structured collection of treatment data as FHIR data (D-7.7)	C †	Trade blog, no interest declared; content traces to the company	Open, volatile
nursIT Institute, reply to an enquiry by the author, 17 August 2026	Personal communication	That careIT writes clinical content directly as FHIR resources through the FHIR client, with no intermediate clinical model, and that the repository is interchangeable where it supports FHIR R4 (D-7.8)	D †	The vendor of the product, on its own architecture, in answer to a question about that architecture. It does not address SENSDOC, names no version and no installation beyond a count	Unpublished, held by the author
nursIT Institute, product pages (careIT, SENSDOC)	Artefact inspected	That careIT captures at the point of care, that Smile CDR is carried as a FHIR-native repository, and that SENSDOC names a database of its own with FHIR for inter-component communication (D-7.8)	D †	The vendor of the products described	Open, volatile

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
van Holt et al. 2026, whitepaper (DMI · ID Berlin · Firemetrics) — no public address established; a search on 13 August 2026 reached only the associated event page, `id-berlin.de/cdr-konferenz`, and not the document itself	Full text, from the copy held by the author	That a second German product marketed as "FHIR-native" describes an Option B architecture in its own words, and the existence of a second integrated commercial CDR offering with named German deployments (§ 1.6.1, D-7.9)	D	Authored by the three supplying companies; the architecture statement runs against the product's own name, the openEHR statement serves the supplier's interest and is used for nothing	Not located publicly; copy held by the author

G-7: Counts Made for This Assessment

The three rows below are not sources but figures the author produced from the counted German package corpus. They stand in this ledger because a grade attaches to them, and the queries and inventories behind them are registered at Appendix H-3 and H-4.

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
German FHIR profile corpus, binding-strength count	Counted	The distribution of binding strengths and fixed-value constraints across the counted corpus, and the safety-critical subset (D-10.7, § 2.5.2)	A — a direct reading of identified artefacts, with the counting rule, the exclusion, the tie-break sensitivity and a digest of the artefact set stated in D-10.7	Direct reading of the artefacts supplied to the author	Restricted; the corpus is not published as such, the count is reproducible from the stated method
German FHIR	Counted	The profile hierarchy of the	A for the count	Author-derived from the same	Not published with the report;

SOURCE	BASIS OF ASSESSMENT	WHAT IT CARRIES	GRADE	INTERESTS	RETRIEVAL
profile corpus, derivation count		corpus: 1,167 of 2,317 resource profiles constrain an HL7 core resource directly, 1,110 derive from another profile in the corpus, 721 of those stay inside one package, 62 name a national base profile as parent (§ 2.6.1; Document 5, D-10.7). The corresponding figures published at release 1.7.0 were wrong, because a version-suffixed parent was read as lying outside the corpus; the correction is recorded in Document 5, D-10.7	against the named artefact set, which reproduces the published corpus totals; no grade for any proposition about semantic reuse	third-party packages as the binding count; the scan is described in Document 5	reproducible from the named packages
German FHIR profile corpus, version-binding count	Counted	That 204 of 2,996 `baseDefinition` references and 232 of 2,379 differential value-set references carry a version, while 325 of 364 package-manifest dependency statements name an exact version; and that of the in-corpus derivations, none of the 721 intra-package ones names a version while 193 of the 389 cross-package ones do (§ 1.5; Document 5, D-10.7)	A for the count against the named artefact set; no grade for any proposition about resolution behaviour in a given tool	Author-derived from the same third-party packages as the binding and derivation counts	Not published with the report; reproducible from the named packages

The rows that carry no persistent identifier, named rather than left to be noticed. A DOI or canonical address was established for every peer-reviewed publication in this ledger. Six rows have none, and each is a different kind of gap rather than an oversight:

- **Three DevDays presentations** (Namlı 2025, Manzelli 2025, Chai 2024) and **one personal communication** (Smithies 2026). Conference slides and recordings are not published in a citable form, and a communication is not a document. A programme or session page would give a reader something to click and would not give them the source, so none is offered here. All four are graded D and are the sole support for no proposition.
- **The Gartner research note of April 2017**, which is behind a subscription and was read as summarised rather than as published. The row already says "reported only" and the grade already reflects it.
- **The van Holt whitepaper**, whose public location could not be established; the search made and what it reached are recorded in the row itself.

The bvitg, DKG and ZVEI consultation statements are consultation documents rather than publications, but the ministry's consultation mirror gives them an address, recorded in the bibliography; the Beale forum post and the HL7 UK article carry their titles and addresses there as well.

What the ledger does not do. It does not raise the evidentiary value of any source. A row recording "Full text" for a vendor white paper leaves it a vendor white paper. The ledger records how the assessment of a source was arrived at, not how much the source can carry; the latter remains the business of the grade and of the appendix in which the source is discussed.

Bibliography

Appendix G records, for the load-bearing entries below, what the assessment of each source was based on, what it carries, its grade, the interests bearing on it and its retrieval status. The summary extract that formed the working basis for part of this literature accompanies the document as Literature Summaries.md.

Working materials (unpublished):

- *Expertendiskussion openEHR* (February 2024), unpublished.
- *Technische Fakten zu openEHR*, unpublished.
- Oemig F, Blobel B (2020): *HL7 FHIR, openEHR, oder was sonst noch?* Internal comparison document, 17 Aug 2020.
- Oemig F: *Comparison of HL7 Product Families.* Excel comparison table, undated.

Published professional-society and conference materials:

- Oemig F, Blobel B, Helmer A, Birkle M (2018): *Vergleich verschiedener Lösungsansätze für Interoperabilität in der Medizin. HL7-Mitteilungen* Nr. 41/2018, pp. 6–17. **ISSN 1616-8909** — published periodical of HL7 Deutschland; public archive at hl7.de/downloads/hl7-mitteilungen (individual issues as PDF). Editorially edited, not peer-reviewed; this is one of the FHIR-leaning source materials characterised in "On methodology and sources".
- Blobel B, Oemig F (2018): *Was läuft schief bei der Modellierung von Interoperabilität? HL7-Mitteilungen* Nr. 41/2018, pp. 18–24. ISSN 1616-8909; same periodical and status as above.
- Pazos P, Atalag K, Marco-Ruiz L, Sundvall E, Miranda Freire S (2015): *Design and Implementation of Clinical Databases with openEHR.* Tutorial presentation, 77 slides. Publicly available — deposited on Figshare and archived on the openEHR resources wiki; also on SlideShare (Pablo Pazos). An openEHR-internal source **by origin** (the "internal openEHR source" note in § 2.2 and Appendix B refers to its camp provenance, not to non-publication).

Peer-reviewed literature and external analyses:

- OECD (2026): *Interoperability in Healthcare: Towards an Interconnected Future.* OECD Health Working Papers No. 197, OECD Publishing, Paris. DOI: 10.1787/eb71b7be-en.
- Piera-Jiménez J, Cano I, Carot-Sans G, Cresswell K, Dunscombe R, Freriks G, Frid S, Hovenga E, Kalra D, Koch S, Leslie H, et al. (2026): *Recommendations for a national electronic health record in Spain: a Delphi study.* Lancet Digit Health. DOI: 10.1016/j.landig.2026.101037. Open access, CC BY-NC. Ministry-commissioned Delphi, 45 items, 152 of 220 experts, single-round consensus; appraised at Appendix E-12. Evidence grade A for the wording, B for the consensus figures, no grade for the suitability verdict.
- Pedrera-Jiménez M, García-Barrio N, Frid S, Moner D, Boscá-Tomás D, Lozano-Rubí R, Kalra D, Beale T, Muñoz-Carrero A, Serrano-Balazote P (2023): *Can OpenEHR, ISO 13606, and HL7 FHIR Work Together? An Agnostic Approach.* J Med Internet Res 25:e48702. DOI: 10.2196/48702.
- Christensen B, Ellingsen G (2016): *Evaluating Model-Driven Development for large-scale EHRs through the openEHR approach.* Int J Med Inf 89, pp. 43–54. DOI: 10.1016/j.ijmedinf.2016.02.004.
- Gartner Inc. (2017): *Evaluate openEHR Standards for Managing Clinical Content Across the Care Continuum.* Research Note, April 2017; the surviving summary attributes it to Jones M and Hallenberg H under the number G00317087. Cited independently by title and date in European Commission, SWD(2018) 126 final of 25 April 2018, footnote 51. Read in summary form only; the sourcing boundary is stated at Appendix D-6, and nothing beyond the summary is attributed to the report.
- Bruthans J, Jiráková K (2023): *The Current State and Usage of European Electronic Cross-border Health Services (eHDSI).* J Med Syst 47, 22. DOI: 10.1007/s10916-023-01920-9.
- Bernstam EV, Warner JL, Krauss JC, Ambinder E, Rubinstein WS, Komatsoulis G, Miller RS, Chen JL (2022): *Quantitating and assessing interoperability between electronic health records.* J Am Med Inform Assoc 29(5):753–760. DOI: 10.1093/jamia/ocab289. Measures a string-recognition interoperability score over 12 structured elements at 68 US oncology sites: 0.68 intra-vendor, 0.22 inter-vendor. Evidence grade B for the scores; the limitations are set out at Appendix D-11.3.
- D'Amore JD, Mandel JC, Kreda DA, Swain A, Koromia GA, Sundareswaran S, Alschuler L, Dolin RH, Mandl KD, Kohane IS, Ramoni RB (2014): *Are Meaningful Use Stage 2 certified EHRs ready for interoperability? Findings from the SMART C-CDA Collaborative.* J Am Med Inform Assoc 21(6):1060–1068. DOI: 10.1136/amiajnl-2014-002883. Open access. Documents semantic errors and permitted variation in C-CDA output. The 615 observations come from 21 manually inspected documents, selected from 91 documents produced by 21 technologies. Evidence grade B for those observations; the limitations are set out in Appendix D-11.1.
- Wright A, Aaron S, Seger DL, Samal L, Schiff GD, Bates DW (2018): *Reduced Effectiveness of Interruptive Drug-Drug Interaction Alerts after Conversion to a Commercial Electronic Health Record.* J Gen Intern Med 33(11):1868–1876. DOI: 10.1007/s11606-018-4415-9. Open access. Supports the narrower § 1.6.2 finding that data migration does not ensure equivalent safety-related application behaviour. Evidence grade A for the measurements at the site studied, with no inference to patient harm or architectural comparison. The legacy system's 100% acceptance figure reflects a hard stop and must not be quoted without that qualification (Appendix D-11.2).
- van Holt J, Nensa F, Sander A, Engels D, Müller A, Kaatze T (2026): *Datensouveränität im Krankenhaus: FHIR-basierte Clinical Data Repositories in der Praxis.* Whitepaper, DMI · ID Berlin · Firemetrics, 27 pages. Vendor documentation for the product Phoenix. Evidence grade D. Carried for one proposition only: the product describes an Option B architecture in its own words while being named FHIR-native (Appendix D-7.9). Its benefit claims, its unsourced opening statistics and its statement about openEHR are expressly not drawn on.
- Huang C, Koppel R, McGreevey JD III, Craven CK, Schreiber R (2020): *Transitions from One Electronic Health Record to Another: Challenges, Pitfalls, and Recommendations.* Appl Clin Inform 11:742–754. DOI: 10.1055/s-0040-1718535. Read in full and **deliberately not drawn on for any proposition**, therefore carrying no ledger row. A narrative review whose own contribution is author consensus, whose cost figures rest on trade press, and whose unit of analysis is the replacement of a whole system rather than the operating path this document

examines. It is listed because it is how D'Amore and Wright were found, and provenance affects selection (Appendix D-11).

- Vorisek CN et al. (2022): *Fast Healthcare Interoperability Resources (FHIR) for Interoperability in Health Research: Systematic Review.* JMIR Med Inform 10(7):e35724. DOI: 10.2196/35724.
- Heyder R et al. (2023): *Das Netzwerk Universitätsmedizin — Technisch-organisatorische Ansätze für Forschungsdatenplattformen.* Bundesgesundheitsblatt 66:108–118. DOI: 10.1007/s00103-022-03649-1. Open access. Appraisal: the architecture report for the NUM research data infrastructure and the source this document cites for the NUM case. It describes the routine-data platform through the GECCO dataset and its FHIR specification and through federated FHIR repositories at the data integration centres; openEHR, ehrbase and the FHIR bridge are not mentioned in it. It therefore carries the FHIR half of the constellation and not the openEHR half (Appendix C-3).
- Sass J, Bartschke A, Lehne M, Essenwanger A, Rinaldi E, Rudolph S, Heitmann KU, Vehreschild JJ, von Kalle C, Thun S (2020): *The German Corona Consensus Dataset (GECCO): a standardized dataset for COVID-19 research in university medicine and beyond.* BMC Med Inform Decis Mak 20:341. DOI: 10.1186/s12911-020-01374-w. The present author is a co-author. Cited in Appendix C-3 only for the paper's own record that the FHIR profiles of the Medical Informatics Initiative were taken into account in defining the GECCO profiles, which dates the direction of influence.
- Ladas N, Franz S, Haarbrandt B, Sommer KK, Kohler S, Ballout S, Fiebeck J, Marschollek M, Gietzelt M (2022): *openEHR-to-FHIR: Converting openEHR Compositions to Fast Healthcare Interoperability Resources (FHIR) for the German Corona Consensus Dataset (GECCO).* Stud Health Technol Inform 289:485–486. DOI: 10.3233/SHTI210963. Short conference contribution; evidence grade B for the narrow statement that COVID data held in openEHR repositories had to be converted into GECCO FHIR for exchange (Appendix C-3).
- Behrend P, Ingenerf J (2022): *Generating GECCO instance data in HL7 FHIR and openEHR using Synthea.* 67th Annual Meeting of the GMDS, 21–25 August 2022. DOI: 10.3205/22gmds069. Conference abstract; evidence grade C. Cited in Appendix C-3 for the statement that profilings of GECCO exist in both HL7 FHIR R4 and openEHR and that both were held equally applicable for the use case.
- Medical Informatics Initiative: *Der Kerndatensatz der Medizininformatik-Initiative* (initiative's own presentation of the core data set), and Ammon D (7 December 2018): *MI-I-Kerndatensatz, ART-DECOR, HL7 FHIR — Technische Optionen für die Governance der Fortschreibung des Kerndatensatzes*, presentation of the MII working group on interoperability at the HL7 Deutschland Interoperabilitätsforum, eleven slides, held on file as `material/HL7DE-MII-Kerndatensatz-2018-12-07.pdf`. Self-reporting by the specifying bodies; evidence grade B. Cited in Appendix C-3 for the adoption of the textual core data set in March 2017 and for the state of the FHIR work on the base modules in December 2018 (mappathon of 9 November 2018; six named profiles as phase-one deliverables); the formal FHIR decision of April 2019 is carried by the *Manteldokument* listed under the regulatory and institutional sources.
- Hyppönen H, Lumme S, Reponen J, Vänskä J, Kaipio J, Heponiemi T, Lääveri T (2019): *Health information exchange in Finland: Usage of different access types and predictors of paper use.* Int J Med Inf 122:1–6. DOI: 10.1016/j.ijmedinf.2018.11.005. National physician surveys of 2010, 2014 and 2017; carries the Finnish usage row of Appendix D-4. Evidence grade C (repeated cross-sectional survey, self-reported use, no outcome measurement).
- Piera-Jiménez J, Carot-Sans G (2025): *Catalonia's Journey to the Open Platform Paradigm in Healthcare.* In: Scheuer A, Studzinski J (eds.), *Digital Maturity in Hospitals.* Springer, Cham. DOI: 10.1007/978-3-031-80704-6_8.
- Allwell-Brown E (2016), **not drawn on for any proposition in this set** and listed for completeness of the search record: *A Comparative Analysis of HL7 FHIR and openEHR for Electronic Aggregation, Exchange and Reuse of Patient Data in Acute Care.* Master's thesis, Karolinska Institutet / Stockholm University. Appraisal: a qualitative-descriptive single case without peer review, technologically based on FHIR DSTU 2 / openEHR 1.0.3 (2015) and

thus around ten years old; individual structural observations are consistent with more recent sources, but citable only with an explicit age caveat.

- Borys K, Hartmann EM, Idrissi-Yaghir A, Livingstone E, Lodde G, Schmidt CS, Winnekens P, Friedrich CM, Hosch R, Nensa F (2025): *DermaDashboard: Bridging the Gap Between FHIR Standards and Clinical Usability.* JMIR Cancer 11:e73691. DOI: 10.2196/73691. Peer-reviewed viewpoint; evidence grade C (Appendix D-7.3).
- Schmidt CS, Wutzkowsky J, Schäfer H, Reinartz L, Chahin H, Winnekens P, Alsulaiman JA, Temme C, Lenz V, Hosch R, Nensa F, Friedrich CM, Böckmann B, Horn PA (2026): *A real-time dashboard for optimizing blood product utilization through a FHIR-based data integration pipeline.* Transfusion and Apheresis Science 65:104301. DOI: 10.1016/j.transci.2025.104301 (online since 19 Dec 2025). Peer-reviewed short report; evidence grade C (Appendix D-7.3).
- Ammon D, Kurscheidt M, Buckow K, Kirsten T, Löbe M, Meineke F, Prasser F, Saß J, Sax U, Stäubert S, Thun S, Wettstein R, Wiedekopf JP, Wodke JAH, Boeker M, Ganslandt T (2024): *Arbeitsgruppe Interoperabilität: Kerndatensatz und Informationssysteme für Integration und Austausch von Daten in der Medizininformatik-Initiative.* Bundesgesundheitsblatt 67(6):656–667. DOI: 10.1007/s00103-024-03888-4. Evidence grade B (Appendix D-7.4). Interests: authored by the initiative's own interoperability working group.
- Kohler S, Piera-Jiménez J, Anywar M, Fuhrmann L, Leslie H, Meixner M et al. (2025): *FHIRconnect: Towards a seamless integration of openEHR and FHIR.* arXiv:2511.14618. DOI: 10.48550/arXiv.2511.14618. Preprint; relevant to the mapping discussion (§ 1.6), not yet peer-reviewed. Its reference engine is **openFHIR** (open-fhir.com), a dockerised, bidirectional openEHR↔FHIR mapping engine offered in open-source and commercial editions. It supplies tooling for the cross-standard boundary in Options D and E; it does not remove that boundary. Evidence grade D for the tool claims, with no independent evaluation.
- Kohler S, Boscá D, Kärcher F, Haarbrandt B, Prinz M, Marschollek M, Eils R (2023): *Eos and OMOCL: Towards a seamless integration of openEHR records into the OMOP Common Data Model.* J Biomed Inform 144:104437. DOI: 10.1016/j.jbi.2023.104437. Open access. Primary source for the openEHR-to-OMOP path in § 2.7: a declarative archetype-to-CDM mapping language (OMOCL) and the Eos tool that executes it on real-world and sample records. Load-bearing for the statement that the projection into the less expressive OMOP CDM loses semantics, which is the authors' own conclusion. Evidence grade B; not to be confused with Kohler et al. 2025, which concerns openEHR-to-FHIR mapping.
- Lenert LA, Ilatovskiy AV, Agnew J, Rudisill P, Jacobs J, Weatherston D, Deans KR (2021): *Automated production of research data marts from a canonical fast healthcare interoperability resource data repository: applications to COVID-19 research.* J Am Med Inform Assoc 28(8):1605–1611. DOI: 10.1093/jamia/ocab108. Evidences that a repository holding clinical data as FHIR canonically can produce PCORnet and OMOP research marts automatically (§ 2.7). Evidence grade C; two authors are employed by a FHIR server vendor, disclosed at the point of use.
- Essaid S, Andre J, Brooks IM, Hohman KH, Hull M, Jackson SL, Kahn MG, Kraus EM, Mandadi N, Martinez AK, Mui JY, Zambarano B, Soares A (2024): *MENDS-on-FHIR: leveraging the OMOP common data model and FHIR standards for national chronic disease surveillance.* JAMIA Open 7(2):ooae045. DOI: 10.1093/jamiaopen/ooae045. Open access. Evidences the reverse leg, OMOP to US-Core-conformant FHIR with Bulk FHIR export, at a non-conformance rate below one percent (§ 2.7). Evidence grade C. The resource count stated in the abstract is implausible by two orders of magnitude and is not relied on.
- Jayathissa P, Rohatsch L, Sauermann S, Hussein R (2025): *OMOP-on-FHIR: Integrating the Clinical Data Through FHIR Bundle to OMOP CDM.* Stud Health Technol Inform 327:667–671. DOI: 10.3233/SHTI250432. Short conference contribution describing the ETL from FHIR Bundles into the OMOP CDM (§ 2.7). Evidence grade D.
- Finster M, Wenzel M, Taghizadeh E (2025): *Common data models and data standards for tabular health data: a systematic review.* BMC Med Inform Decis Mak 25(1):422. DOI: 10.1186/s12911-025-03267-2. Open access. Compares i2b2, the Sentinel CDM, the PCORnet CDM and OMOP against CDA, HL7 v2, FHIR and openEHR across suitability, popularity, adaptability, interoperability and support; OMOP and FHIR score highest overall and the

review concludes that no single global representation can be selected (§ 2.7). Evidence grade B; the criteria include popularity and support, and the scope is tabular health data.

Systematic reviews, and the implementation studies identified through them (Appendix D-8): Three of the rows below are secondary studies, Vorisek, Tabari and Ayaz; Duda (D-8.5) and Finster (§ 2.7) are listed with the groups they are cited from. The remainder are primary implementation reports that those reviews surfaced and that are cited here in their own right, not as reviews.

- Vorisek CN, Lehne M, Klopfenstein SAI, Mayer PJ, Bartschke A, Haese T, Thun S (2022): *Fast Healthcare Interoperability Resources (FHIR) for Interoperability in Health Research: Systematic Review.* JMIR Med Inform 10(7):e35724. DOI: 10.2196/35724. PROSPERO CRD42021235393. **Evidence grade B**; methodologically the strongest of the three (pre-registered). Publication bias expressly conceded by the authors. Interests: the last author is deputy chair of HL7 Germany (declared).
- Tabari P, Costagliola G, De Rosa M, Boeker M (2024): *State-of-the-Art FHIR-Based Data Model and Structure Implementations: Systematic Scoping Review.* JMIR Med Inform 12:e58445. DOI: 10.2196/58445. Evidence grade B/C; a scoping review without systematic quality assessment.
- Ayaz M, Pasha MF, Alzahrani MY, Budiarto R, Stiawan D (2021): *The FHIR Standard: Systematic Literature Review of Implementations, Applications, Challenges and Opportunities.* JMIR Med Inform 9(7):e21929. DOI: 10.2196/21929 (corrected version; correction: 10.2196/32869). Evidence grade C; purely structural quality assessment, no exclusions. Usable as an inventory, not as an argument.
- Rinaldi E, Thun S (2021), **two publications of the same work on the HiGHmed infection-control use case**: *From openEHR to FHIR and OMOP data model for microbiology findings*, Stud Health Technol Inform 281:402–406, DOI: 10.3233/SHTI210189; and *Mapping from openEHR to FHIR and OMOP CDM to support interoperability for infection control*, GMS Med Inform Biom Epidemiol 17(2):Doc07, DOI: 10.3205/mibe000221, eight pages with an attachment giving the element-by-element map and its equivalence scores. A documented openEHR → FHIR → OMOP transformation path; load-bearing for § 1.6 and § 2.4.1, and for § 2.7 and § 2.6 through the journal article. **Neither quantifies effort**; the journal article does score equivalence against ISO/TR 12300, which the conference paper does not. Interests: the second author is deputy chair of HL7 Germany and chairs the Interop Council appraised at Appendix E-11, the first author is a member of her group at the Charité, and the Charité belongs to HiGHmed — the assessment of the mapping is therefore made by the implementing party about its own work, and is used as such. Evidence grade C.
- Hund H, Wettstein R, Heidt CM, Fegeler C (2021): *Executing distributed healthcare and research processes — the HiGHmed data sharing framework.* Stud Health Technol Inform 278:126–133. DOI: 10.3233/SHTI210060.
- Wettstein R, Kussel T, Hund H, Fegeler C, Dugas M, Heidt CM (2021): *Feasibility queries in distributed architectures — concept and implementation in HiGHmed.* Stud Health Technol Inform 278:134–141. DOI: 10.3233/SHTI210061.
- Wettstein R, Hund H, Fegeler C, Heidt CM (2021): *Data sharing in distributed architectures — concept and implementation in HiGHmed.* Stud Health Technol Inform 283:111–118. DOI: 10.3233/SHTI210548. The three contributions together demonstrate that the HiGHmed consortium, which is openEHR-based, operates its cross-site data sharing FHIR-based; used in § 1.6.1. The author lists are as recorded in the review supplied on 13 August 2026 and were not verified against the publisher pages, which is why the DOI rather than the author list is the identifier to follow.

Normative specification sources (load-bearing for § 2.5.1 and § 1.3 Dimension 4):

- openEHR International: *Data Types Information Model*. Reference Model, latest release. Normative for the statement that `terminology_id = "local"` means the origin of the permissible values is described within the archetype. Evidence grade A (specification).
- openEHR International: *Architecture Overview, ch. 12 — Terminology in openEHR*. Normative for the double function of local at-codes (node identifiers and values), for the definition of the archetype path as an

alternating pattern of RM attribute names and node codes, and for the self-stated caveat that coverage by external terminologies is usually incomplete and the mappings approximate. Evidence grade A.

- openEHR International: *Support Terminology Specification.* — Complementary, for the delimitation of the openEHR-internal terminology from external terminologies.
- freshEHR / Digital Health Innovation Scotland: *openEHR Datatypes* and *Understanding composition Constraints and validation.* — Vendor-affiliated implementation documentation; evidences that in the simplified (FLAT) format value and terminology are dispensable because predefined in the template, and that the operational template (`.opt`) is regarded as the authoritative validation source. Evidence grade C; relevant as an **admission against interest**, coming from the openEHR-affiliated environment.
- HL7 International: *FHIR Data Types* — `Coding`, `CodeableConcept`. Specification, R4 and R5. Normative for the statement that `Coding.system` is `0..1` and that a `CodeableConcept` carrying `text` alone is valid, and therefore for the delimitation in § 2.5.1 that the globally resolvable pointer is the normal case rather than a guarantee of the data type. Evidence grade A (specification).
- HL7 International / Clinical Interoperability Council: *FHIR R4 Symptoms Implementation Guide*, STU 1 (build April 2026), section "Logical Model, OpenEHR Archetype, and FHIR Profile". Primary source for Appendix E-7; evidences that the FHIR Logical Model retains the authoritative status and the openEHR archetype is carried as an *informative* document.
- HL7 International: *Structured Data Capture (SDC) Implementation Guide*, continuous build, retrieved 11 August 2026. The extension set and conformance rules on `Questionnaire` and `QuestionnaireResponse` that the comparison in § 2.6.1 (derivation at D-15) turns on, including data extraction into clinical resources by way of `StructureMap`. Evidence grade A for what the guide specifies, and no grade for how widely it is implemented.
- HL7 International: *FHIR QuestionnaireResponse, search parameters*, R4 (hl7.org/fhir/R4), R5 and the continuous build (build.fhir.org), retrieved 11 August 2026. Source for the query-capability statement in § 2.6.1 (derivation at D-15): R4 defines no parameter on item answers or `linkId`; R5 adds `item-subject` only; the current build adds composite `item-*` parameters pairing `linkId` with the answer value. Evidence grade A for what each release defines.
- HL7 International: *Structured Data Capture — Modular Forms* and *Assemble Modular Questionnaire Operation*, retrieved 11 August 2026. Source for the composition mechanism described in § 2.6.1 (derivation at D-15): the `subQuestionnaire` extension references another Questionnaire by canonical URL, which the guide recommends be version-specific, and `$assemble` expands the references recursively into a single form. The guide records that research organisations maintain question libraries and ships one as an example artefact; it does not establish a curated library itself. Evidence grade A for what the guide specifies.
- AKTIN — Nationales Notaufnahmeregister: *AKTIN FHIR Implementation Guide*, publication of 31 July 2025, in particular the technical background, retrieved 11 August 2026; project site aktin.org. Primary source for Appendix D-13 and for the observation in § 2.6.1 (derivation at D-15) that register reporting rests on derived facts: Questionnaires pre-populated from an AKTIN Record built out of `Observation`, `ServiceRequest` and `MedicationAdministration`, the stated principle of no shadow data management beside care, derived age granularity as a computation over source dates, and annotation with LOINC, SNOMED CT and ATC in preference to closed answer lists. Evidence grade B for what the guide specifies. Published through the ART-DECOR environment, in which the author is involved; disclosed at D-13.
- REDCap: *Shared Library*, Vanderbilt University Medical Center, retrieved 11 August 2026. The real-world counter-example in § 2.6.1 (derivation at D-15): a curated holding of data-collection instruments, reviewed by the REDCap Library Oversight Committee for research relevance, correct functioning and copyright before inclusion, and importable into projects. Access is restricted to authenticated users at partner institutions, so the holding itself was not inspected; evidence grade C for the existence and the described curation procedure.
- Xt-EHR / European Commission: *EHDS Logical Information Models*, implementation guide. Second primary source for Appendix E-7; evidences that logical models are rendered by the ordinary FHIR publication toolchain as

hierarchical structure tables (snapshot, differential, key-elements views) with cardinalities, data types, bindings and obligations, and that the UML class diagrams it contains sit at use-case level and are authored alongside the models rather than generated from them. Published guide, verifiable by inspection; evidence grade B.

- ART-DECOR Open Tools: *DECOR datasets, transactions and the version-aware FHIR endpoints* (docs.art-decor.org, product documentation). Source for the ART-DECOR observation in Appendix E-7, including the export asymmetry by which a dataset carries no cardinality and exports with `0..*` while a transaction transfers its cardinality into the logical model. Tool documentation without independent evaluation; **evidence grade D**. The author of this assessment is involved in the ART-DECOR Expert Group; the interest is disclosed in Appendix E-7 itself.
- HL7 Europe / IHE-Europe: *EURIDICE — European Interoperability Specifications for Digital Solutions in Healthcare*, specification catalogue, and *HL7 Europe and IHE-Europe open coordinated ballots for European FHIR Implementation Guides developed jointly in the EURIDICE collaboration*, retrieved 11 August 2026. Primary source for Appendix D-12 and for the specification mechanism described in § 1.1: the catalogue of European FHIR implementation guides (EU Patient Summary, laboratory report, hospital discharge report, imaging report, MADO, medication prescription and dispense, EU Health Data API) and the coordinated ballot of the Imaging Study Report IG and the EU Health Data API IG with a deadline of 30 April 2026. Published artefacts read for what they state; **evidence grade B**. The author holds roles in HL7 Deutschland and HL7 Europe; the interest is disclosed in Appendix D-12 itself.
- openEHR International / HL7 International (May 2026): *Converge & Collaborate 2026 — Joint Statement*. Press statement on the second joint annual meeting, Dublin, 12/13 May 2026. Institutional self-presentation; see Appendix E-8.

Position, framework and further analytical documents (Appendices A, D-8, D-14, E and § 1.2):

- Tsafnat G, Dunscombe R, Gabriel D, Grieve G, Reich C (2024): *Converge or Collide? Making Sense of a Plethora of Open Data Standards in Health Care.* J Med Internet Res 26:e55779. DOI: 10.2196/55779. **Editorial, not peer-reviewed**; the authors are institutional representatives of the standards recommended; conflicts of interest declared. Evidence grade D. See Appendix E-1.
- Kapitan D et al. (2025): *Data Interoperability in Context: The Importance of Open-Source Implementations When Choosing Open Standards.* J Med Internet Res (Viewpoint). DOI: 10.2196/66616. Peer-reviewed counter-position to Tsafnat et al.
- Schabetsberger T, Mangesius P (2026): *HL7 FHIR and openEHR: A Comparative Analysis, Implementation Strategies, and Recommendations for Healthcare Organizations and Regions.* Technical White Paper v1.0, Siemens Healthcare Diagnostics GmbH, Vienna, June 2026. **Vendor paper, self-declared not peer-reviewed.** Evidence grade C/D. See Appendix E-2.
- ÖGTelemed (2026): *Positionspapier zu herstellerneutralen, applikationsunabhängigen Gesundheitsdaten.* Version 1.0, 24 June 2026. Position paper by a **specialist society** (not the medical chamber or a public authority). Evidence grade C. See Appendix E-3.
- Australian Digital Health Agency (2026): *National Framework for Digital Health Standards*. ISBN 978-1-7643527-2-7, May 2026. Official national framework. Evidence grade B for governance statements. See Appendix E-4.
- Fuhrmann L et al. (2026): *Developing a Vendor-Neutral, Context-Inclusive Clinical Information Model for Intraocular Pressure.* Conference poster P.178, *openEHR eye care collab*. Evidence grade D; evidences the CKM governance process, not architecture. See Appendix E-5.
- Devitt D (28 May 2026): *Is openEHR the Linux of clinical data standards?* darrendevitt.com. Professional blog of a single author, no peer review, no source references. Evidence grade D. The statement runs against the apparent interest of its author, an openEHR-leaning writer reaching a market conclusion the other way. That lowers the likelihood of simple partisanship and is recorded as such, but it is **not** treated as evidential weight

and does not raise the grade: a writer can argue against an apparent camp for strategic, rhetorical or personal reasons, and the assignment of an author to a camp is itself partly interpretive. Supplies, with the "Red Hat problem", a distinct lock-in mechanism (§ 1.3, Dimension 1). His premise of openEHR's technical superiority is expressly **not** adopted. See Appendix E-6.

- Devitt D (2026): *FHIR server providers* (market catalogue, updated July 2026). darrendevitt.com. An informal but current inventory listing on the order of thirty FHIR server products (hyperscaler, commercial and open-source). Evidence grade D; used in § 1.3 (Dimension 1) and Appendix A-1 as a market-breadth data point.
- Delussu, G., Frexia, F., Mascia, C., Sulis, A., Meloni, V., Del Rio, M., Lianas, L. (2024): *A survey of openEHR Clinical Data Repositories.* International Journal of Medical Informatics 191, 105591. doi:10.1016/j.ijmedinf.2024.105591. Peer-reviewed survey of eleven of seventeen identified openEHR repositories, with the answer set and analysis code published for reuse. Evidence grade B for the declared product properties, C for inference to deployment behaviour. See Appendix A-2.
- Kramer MA (2023): *Reducing FHIR "Profiliferation": A Data-Driven Approach.* AMIA Annu Symp Proc 2022:634–643. PMID 37128432, PMC10148270. A review of 2,897 profiles from 125 implementation guides, finding that two profiled items, Observation and Extension, account for 49% of the profiles sampled, that standard extensions make up only about one fifth of the extensions used, and that FHIR's 80/20 rule is closer to 65/35, with profiling defining 33 to 38% of new data elements. Full text read. Cited in Appendix E-10 for what it does and does not establish about the OECD's paragraph 20. Not to be confused with Kramer & Moesel 2023, a different paper by one of the same authors. Evidence grade B.
- Kramer MA, Moesel C (2023): *Interoperability with multiple Fast Healthcare Interoperability Resources (FHIR) profiles and versions.* JAMIA Open 6(1), ooad001. DOI: 10.1093/jamiaopen/ooad001. A methods contribution proposing the FHIR Interoperability Table, with catchment and admittance as its two concepts, demonstrated on US Core 5.0.1 against the International Patient Summary 1.0.0. Evidence grade A for the computed relationships, B for the method as a general instrument. See Appendix D-14.
- Jones J, Gottlieb D, Mandel JC, Ignatov V, Ellis A, Kubick W, Mandl KD (2021): *A landscape survey of planned SMART/HL7 bulk FHIR data access API implementations and tools.* J Am Med Inform Assoc 28(6):1284–1287. DOI: 10.1093/jamia/ocab028. Evidence grade C; 22 of 37 organisations approached, self-reported plans, United States, 2020. Used in § 1.2, point 2.
- Mandl KD, Gottlieb D, Mandel JC, Ignatov V, Sayeed R, Grieve G, Jones J, Ellis A, Culbertson A (2020): *Push Button Population Health: The SMART/HL7 FHIR Bulk Data Access Application Programming Interface.* npj Digital Medicine 3:151. DOI: 10.1038/s41746-020-00358-4. Architecture and advocacy paper; evidence grade C. Used in § 1.6 for its account of what custom site-specific mapping costs, which is a statement against the authors' rhetorical position rather than for it.
- Duda SN, Kennedy N, Conway D, Cheng AC, Nguyen V, Zayas-Cabán T, Harris PA (2022): *HL7 FHIR-based tools and initiatives to support clinical research: a scoping review.* J Am Med Inform Assoc 29(9):1642–1653. DOI: 10.1093/jamia/ocac105. Evidence grade B/C; 203 projects from 1,572 candidates. See Appendix D-8.5.
- Griffin AC, He L, Sunjaya AP, King AJ, Khan Z, Nwadiugwu M, Douthit B, Subbian V, Nguyen V, Braunstein M, Jaffe C, Schleyer T (2022): *Clinical, technical, and implementation characteristics of real-world health applications using FHIR.* JAMIA Open 5(4), ooac077. DOI: 10.1093/jamiaopen/ooac077. Brief communication, convenience sample, evidence grade C. Used in Appendix A-1 for the application-layer picture.
- Health Samurai (29 June 2026): *FHIR Server Performance at Scale: Baseline*. Vendor benchmark of Aidbox, HAPI FHIR, Medplum and the Microsoft FHIR Server; harness and results published at github.com/HealthSamurai/fhir-server-performance-benchmark. Evidence grade D; the publisher builds the leading product and the dataset is 1,000 patients. See Appendix A-1.
- CTSA data-harmonization project (Simhadri S, Hemadri R, Chopra V): *FHIR repository evaluation*. Comparison of seven FHIR repository products across seven criteria, referencing HAPI 3.8.0 and 4.0.0 and therefore dating

from 2019. Evidence grade D and stale; the evaluators state they worked from documentation without expert consultation. See Appendix A-1.

- Marfoglia, A., Arcobelli, V. A., Moscato, S., La Mattina, A. A., Mellone, S., Carbonaro, A. (2026): *Challenges of health data standard adoption and usage: a systematic review.* Journal of Biomedical Informatics 178, 105022. doi:10.1016/j.jbi.2026.105022. PROSPERO CRD42024623398. Open access, CC BY.
- openEHR Switzerland / openEHR International (2026): *From proof of concept to a reusable blueprint.* openehr.org (guest post, J.-P. Messerli). Announcement of a Swiss openEHR and HL7 joint working group and a vaccination-showcase blueprint (FHIRConnect/openFHIR, CH-VACD). Camp self-presentation of a hackathon proof of concept; evidence grade D. See Appendix E-8 and Appendix C-1.

Wales / NHS Wales sources (Appendix D-7.2):

- Welsh Government (14 June 2023): *Introduction of HL7 FHIR as a foundational standard in all NHS Wales Bodies*. Welsh Health Circular WHC/2023/018.
- Digital Health and Care Wales: *Wales FHIR Implementation Guide* (v2.6.0), Simplifier.net.
- HL7 UK (4 June 2026): *Welsh FHIR: Progressing interoperability across Wales*. Trade article with Callum Saint (DHCW). Institutional self-presentation, no independent evaluation.

Further FHIR production cases (Appendix D-7.5):

- Namlı T (DevDays 2025): *Common Data Model implementation in Türkiye.* Conference presentation, devdays.com (slides and recording). Source for the reported holding of some 38 billion FHIR resources on around 32 million patients, and for the access times cited there. Presentation by the implementing party, no independent evaluation; evidence grade D.
- Manzelli V (DevDays 2025): *Implementing a FHIR-based regional CDR.* Conference presentation, devdays.com (slides and recording). Source for the Lombardy regional clinical data repository at around 10 million patients. Presentation by the implementing party, no independent evaluation; evidence grade D.
- Chai V (DevDays 2024): *HL7 FHIR Implementation for Singapore National Healthier SG Programme.* Conference presentation, devdays.com (slides and recording). Source for the Singapore case. Presentation by the implementing party, no independent evaluation; the material does not determine the persistence model; evidence grade D.
- Smithies R (HL7 UK): personal communication, July 2026, referring in part to the thread *Large FHIR EHRs* on chat.fhir.org. Source for the Ukraine figures (approximately 39 million patients, 420,000 users) and for the qualification that the REST interface there is not a FHIR API. Unpublished communication without a published primary source, from a source active on the HL7 side; evidence grade D.

Regulatory and institutional sources:

- Medizininformatik-Initiative: *Medizininformatik Initiative — ImplementationGuide — Manteldokument*, Simplifier.net, retrieved 18 August 2026. Records the adoption of the first, textual core dataset in March 2017 and the National Steering Committee decision of April 2019 for HL7 FHIR as the technical representation standard. Evidence grade B; the initiative's own account. See also *Der Kerndatensatz der Medizininformatik-Initiative* and the medication-module repository, github.com/medizininformatik-initiative/kerndatensatzmodul-medikation, state August 2019.
- Federal Office of Public Health / DigiSanté (19 May 2026): *Der Standard HL7 FHIR als Fundament des SwissHDS*; announcement by eHealth Suisse of the same date.
- *Wie wird der Nutzen des E-Rezepts von Versicherten und Leistungserbringenden eingeschätzt?* (21 January 2026): gematik GmbH. Not drawn on for any proposition in this set.
- Ernst & Young AG for the BAG (21 November 2025): *Schlussbericht — Standardisierung im SwissHDS mit FHIR.*

- Federal Ministry of Health (15 July 2026): *Kabinett beschließt Gesetz für Daten und digitale Innovation im Gesundheitswesen (GeDIG)* — press release on the cabinet decision. The draft texts are published as the departmental draft and the cabinet draft. The cabinet draft, 265 pages, dated 14 July 2026, was supplied to the author as a file and read in full on 4 August 2026; the earlier sourcing boundary is lifted and the record of it is kept at Appendix D-10.3. Load-bearing for § 1.6.3 and D-10.3, at evidence grade A for what the draft text says.
- HL7 International: *Using Codes in Resources — binding strength* — normative definition of `required`, `extensible`, `preferred` and `example`. Load-bearing for § 2.5.2, evidence grade A for what the specification requires.
- openEHR International: *Archetype Definition Language (ADL2), § 7.13.4 Term_bindings* — the optional term-binding sub-section of archetypes. Load-bearing for § 2.5.2, evidence grade A.
- Beale T (19 August 2020): *Addition of required / extensible / preferred / example terminology constraint to AOM and ADL* — proposal on behalf of the openEHR Specifications Editorial Committee, naming the HL7 FHIR model as the template; discussion continues into 2023 without visible conclusion. Evidence grade D, weighted higher as a statement against the interest of the camp making it. Load-bearing for § 2.5.2.
- E-HEALTH-COM: *GeDIG-Entwurf — konziliant im Ton, hart in der Sache* — trade report on the draft, evidence grade C. It is retained as background but is no longer load-bearing because the primary draft text has been read directly.
- Consultation statements on the GeDIG departmental draft, May 2026: bvitg, 17 May 2026; ZVEI, 15 May 2026; DKG, 18 May 2026. Evidence grade B for the positions stated.
- gematik / Competence Centre for Interoperability in Health Care: *Konformitätsbewertung*, description of the field of action, and *KOB sorgt mit Standards für Transparenz — 141 Systeme bereits erfolgreich bewertet*, 11 December 2025. Both retrieved 11 August 2026. Source for Document 5, D-10.8 and for the qualification in § 1.7 that the machinery a conformance lever requires is buildable and built nationally. The statutory obligation under §§ 385 and 387 SGB V is grade A for what the statute requires; the published description of the procedure is grade B; the figures of 141 assessed systems and some 98% market coverage are self-reported by the operating body and are carried at **grade D** with their date. They evidence the existence and reach of the procedure, not the clinical usability of what passes it.
- European Commission: *European Health Data Space Regulation* (EU) 2025/327.
- HL7 Europe: *FHIR Implementation Guides to support EHDS.*
- HL7: *SQL on FHIR Specification*; *FHIR Bulk Data Access (Flat FHIR).*
- Medical Informatics Initiative / TMF: *Forschungsdatenportal für Gesundheit (FDPG)*, forschen-fuer-gesundheit.de — public aggregate data-holdings counter of the participating data integration centres (figures as of 22 July 2026: 41 reporting centres; 3+ billion laboratory values; 27+ million persons). Self-reported live aggregate, publicly auditable but not independently verified; evidence grade C. Load-bearing for Appendix D-7.4 and Finding 3.

German ePA, profile families and interoperability governance (Appendix D-10):

- gematik GmbH: *gemSpec_Aktensystem_ePAfueralle*, V1.8.2 (14 July 2026); *gemILF_PS_ePA — Implementierungsleitfaden Primärsysteme ePA*, V3.8.1 (24 March 2026). Normative specifications with version and date. Evidence grade A for what they require a system to do, and no grade for the proposition that the requirement is met in the field. Load-bearing for Appendix D-10.1.
- gematik GmbH: ePA FHIR implementation guides, `gemspec.gematik.de/ig/fhir/epa/1.3.0/general-principles.html` and the migration chapter of `de.gematik.epa.medication` 1.3.2. Source for the server-side handling of `id` and `meta.profile` and for the documented migration rule. Evidence grade A for that prescribed behaviour, which is what is cited from them here.

- FHIR package-registry metadata for `de.basisprofil.r4`, `de.gematik.epa*`, `kbv.mio.*`, `kbv.ita.*`, the MII core-dataset packages and the DEMIS packages: versions, dependencies and release dates. Public registry data, read directly from the published manifests; evidence grade A for the declared versions, dependencies and release dates, consistent with the grading of the dependency findings in Appendix D-10.2. The per-family profile counts in Appendix D-10.2 are the author's own summation across published module counts, not published aggregates, and are marked as such at the point of use.
- Interop Council (Germany): *Positionierung des Interop Council zu den Interoperabilitätsvorgaben im Rahmen des GeDIG*, Berlin, 11 August 2026, two pages, published on the national interoperability platform, `ina.gematik.de/community-hub/community-formate/interop-council/positionierung-zu-gedig`, retrieved 18 August 2026. Read in full. Evidence grade A for its wording, B for the positions it takes. Load-bearing for Appendix E-11 and for D-10.3; cited in § 1.2 and § 1.6.3.
- Interop Council (Germany), core-profile working group: position paper on cross-sector core profiles, v1.0, dated 14 March 2025 in the document-history table and 17 March 2025 on the cover page, published on the national interoperability platform. Self-criticism by the responsible governance body, with named working group, version and date; evidence grade B. Load-bearing for Appendix D-10.3 and for the cross-reference in § 1.5.
- Interop Council (Germany), working group on the adoption of the hospital interfaces: position paper of 22 February 2024, based on 219 responses. Evidence grade B; the adoption percentages within it are survey self-reports.
- gematik: interview in the trade press, 9 April 2024, on the state of the hospital-interface specification. Evidence grade C.
- ePA usage figures: survey of the large sickness funds reported in the press for the turn of 2025/2026, the largest fund's own figure of April 2026, and the weekly retrieval figure for the electronic medication list. Press-reported self-reporting; evidence grade C.
- Commercial-register entry for mio42 GmbH, registered 29 October 2020, for the registered object on which the glossary entry in Document 3, Appendix F-5, rests. Evidence grade B.

Bibliometric data sources (Appendix D-9):

- OpenAlex (api.openalex.org): title/abstract counts for `FHIR` and `openEHR`, grouped by publication year, and the sensitivity checks of Appendices D-9.1 to D-9.5: search-string tiers, document types, countries of affiliation, random samples with a fixed seed, and three denominators for the normalised trend. First query date 23 July 2026 (totals: FHIR 3,421; openEHR 837); full re-run and sensitivity checks 2 August 2026 (totals: FHIR 3,505; openEHR 837), twenty-one queries in all, each recorded as a full URL. Open, reproducible bibliographic database; counts research attention, not effectiveness. Recent years (2025–2026) provisional owing to online-first indexing. The queries were run from a browser rather than from the working environment in which this document is maintained, because the API refused requests from it; the date given is the date of that retrieval.
- PubMed / MEDLINE (pubmed.ncbi.nlm.nih.gov): conservative cross-check via `FHIR[tiab]` and `openEHR[tiab]`, and the review counts of Appendix D-9.3 via `FHIR[tiab] AND review[pt]` and `openEHR[tiab] AND review[pt]`. Query date 23 July 2026 (totals: FHIR 1,055; openEHR 291) and 2 August 2026 (totals: FHIR 1,061; openEHR 291; reviews 94 and 7). Europe PMC (`SRC:MED`) returned 955 for FHIR on 23 July 2026, consistent with the PubMed order of magnitude.

FHIR CDR product sources: HAPI FHIR, Firely Server, Smile CDR, Microsoft Azure Health Data Services, Google Cloud Healthcare API, Aidbox/Health Samurai, Medplum, LinuxForHealth/IBM FHIR Server.

Swiss and German hospital landscape (Appendix C): Press releases and reports by Insel Gruppe, USB Basel/OWT/Swisscom/Better, CISTEC, openEHR Switzerland, as well as Charité, the NUM research data platform, HiGHmed and the Medical Informatics Initiative; media reports from 20min, Der Bund, medinside, Netzwoche, kma-online, heise online, bibliomedmanager (2024–2026).

The individual reports are not listed here one by one, and no proposition in this document rests on that body of press reporting alone. Every statement it supports is also carried by a source listed in Appendix G. It is named as the origin of the background picture and nothing more.

German vendor and programme sources (Appendix D-7.7, D-7.8), all retrieved 5 August 2026: Tiplu GmbH, TipluDB product page, tiplu.de; Tiplu GmbH, interview "HL7 FHIR bei der Tiplu GmbH — von wegen nur kommunizieren", magazin.hl7.de; e-health-com, company news on the BMG-funded project *Entwicklung eines intersektoralen FHIR-Datenmodells*, e-health-com.de; medconweb, report on the same project, medconweb.de; nursIT Institute, product pages nursit.de and nursit.de/sensdoc.

Other: LinkedIn post on the FDMG decision of 19 May 2026, superseded by the primary sources listed above and not drawn on.

Appendix H: Primary Artefact Register

What this register is for, and why it is separate from the ledger. Appendix G records judgements about sources. This register records the artefacts that were read directly and that carry no judgement to record. A statute says what it says, a normative specification prescribes what it prescribes, a package manifest declares what it declares.

A reader checks these by opening them at the location given, not by weighing them, which is why they carry a retrieval point instead of an evidence grade. Separating the two keeps the ledger a register of judgement and stops the audit trail from being incomplete, which is the failure the separation was previously open to.

What a row means. The artefact was read at the location and in the version given, and one or more propositions in the set rest on that reading. A proposition may go beyond what the artefact states, that a prescribed behaviour is achieved in practice for instance. Such a proposition is not carried by the artefact, and it appears in Appendix G with a grade instead.

H-1: Legislation and draft legislation

ARTEFACT	VERSION OR DATE READ	CARRIES
Regulation (EU) 2025/327 (European Health Data Space)	In force since March 2025	The regulatory frame, the Implementing Act deadline of March 2027 (§ 1.1)
SGB V § 355	As amended, read 2026	The statutory reference to FHIR for the record-type specifications, and the KBV mandate (§ 1.1, D-10.3)
SGB V § 373	As amended, read 2026	The gematik mandate for the hospital interfaces (D-10.3)
SGB V § 385, § 387	As amended, read 2026	The Interop Council, the binding interoperability requirements and the conformity-assessment mandate of the competence centre (D-10.3, Document 5)
GeDIG cabinet draft, §§ 386a and 386c SGB V	Cabinet approval 15 July 2026	The exchange obligation, the prohibition on formats preventing complete semantic reconstruction, and the absence of an operative persistence duty (D-10.3, § 1.6.3). Also carried in Appendix G, because the divergence between summary sheet and operative text is a judgement

ARTEFACT	VERSION OR DATE READ	CARRIES
Interop Council, *Positionierung … im Rahmen des GeDIG*	11 August 2026, two pages	The statutory advisory body asking for a duty on semantically standardised data holding at the logical layer, and for FHIR as the leading exchange format (E-11, D-10.3). Also carried in Appendix G, because its positions are graded

H-2: Normative specifications and implementation guides

ARTEFACT	VERSION OR DATE READ	CARRIES
HL7 FHIR, *Overview — Architects*	R4 (v4.0.1), retrieved 18 August 2026	The specification's own statement of purpose: interfaces for interoperating between applications, interoperability as the primary purpose; no statement on persistence (§ 2.3)
HL7 FHIR, resource list	R5 (v5.0.0), retrieved 16 August 2026	The count of 157 resource definitions, stated on the page itself (§ 2.4)
HL7 FHIR, `ElementDefinition.binding.strength`	R4/R4B specification text	The four binding strengths and what each obliges (§ 2.5.2). Also carried in Appendix G, where the proposition drawn from it is graded
HL7 FHIR, search parameters for `QuestionnaireResponse`	R4, R5 and the current build	The absence of an item-answer search parameter in R4, `item-subject` in R5, and the composite parameters in the build (§ 2.6.1 (derivation at D-15)). Also carried in Appendix G, where the proposition drawn from it is graded
HL7 FHIR Structured Data Capture	Published guide	`subQuestionnaire`, the `$assemble` operation, and the absence of a curated module library (§ 2.6.1 (derivation at D-15)). Also carried in Appendix G, where the proposition drawn from it is graded
US Core	5.0.1	The sender side of the interoperability analysis at D-14
International Patient Summary	1.0.0	The receiver side of the same analysis
openEHR Archetype Object Model / ADL, § 7.13.4	Current specification	The optional `term_bindings` section and the absence of a binding strength (§ 2.5.2). Also carried in Appendix G, where the proposition drawn from it is graded
openEHR architecture overview, BASE Release 1.0.2	Read 2026	The separation of EHR and demographic information (A-2), and chapter 12 on terminology (§ 2.5.1). Also carried in Appendix G, where the proposition drawn from it is graded
`de.fhir.medication` and the counted German package set	As at the 12 August 2026 corpus	The dependency declarations and the package manifests (Document 5, D-10.2, D-10.7)

H-3: Registry and package manifests

ARTEFACT	VERSION OR DATE READ	CARRIES
The 46 published package archives and 66 repositories of the counted corpus	Counted 12 August 2026	Every corpus figure in Document 5. The corpus definition, the counting rules and the audit record are in Document 5, D-10.7, and the artefact set is identified there by a digest
gematik ePA medication service dependency listing	Release line 1.3.5, retrieved 2 August 2026	That the published listing includes `de.fhir.medication` 1.0.7, and that it is a resolved tree rather than a direct manifest (Document 5)

H-4: Recorded queries

ARTEFACT	DATE	CARRIES
OpenAlex and PubMed title-and-abstract queries for FHIR and openEHR	First retrieval 23 July 2026, baseline and six sensitivity checks 2 August 2026, twenty-one queries in all	Every figure in Appendix D-9. The queries themselves are recorded in D9-SENSITIVITY-QUERIES.md

H-5: Where web artefacts are recorded

Product pages, catalogues, position papers and announcements inspected directly are listed with their retrieval dates in Appendix G rather than here, because each of them carries an evidence grade: a vendor page and a working-group paper are both artefacts, but reading either involves weighing whose statement it is. The dividing line is not the medium but whether a judgement had to be made.

Method, and the use of a large language model. The literature search, the collation of sources, the drafting and the consistency 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. Three limits of the method belong on the record.

The reading boundary. The peer-reviewed publications were read in full. For much of the remaining literature a summary extract produced by a large language model formed the working basis, and it was that extract rather than the publication that was worked from. The extract also covers a number of peer-reviewed publications, so the boundary does not run along the peer-review line and cannot be described by that shorthand.

Appendix G of this document records, source by source, which basis applies. The extract accompanies the set as Literature Summaries.md, so that a reader can see what the working basis actually said and can check the grades it proposes against the grades carried here.

The boundary is a record in one direction only. Every entry in Appendix G other than "Full text" rests on something that can be checked independently of the author. The extract exists and names its sources, the artefact sits at the cited location, a conference presentation has no full text to read. The full-text entries rest on the author's account: they record that the complete publication rather than a summary was the working basis, not that the study was replicated or its data checked.

A further reservation belongs beside it. A number of ledger rows carry a basis of assessment that was reconstructed analytically from how the source is used in this text, rather than reported by the author. Correcting those against his own recollection is an outstanding task. Both asymmetries are stated rather than smoothed over.

A model does not weigh evidence, but it did propose. For the sources on which the summary extract proposes an evidence grade, the grade this document carries is in each case the grade proposed. Whether that agreement reflects influence or independent arrival at the same judgement cannot be determined after the fact, and no claim is made either way.

At version 3.16.0 every grade carried by a load-bearing source was re-derived from the rubric now stated at the head of Appendix G, without reference to what the extract had proposed. The outcome is recorded in the development log. **That re-derivation is not an independent second assessment and must not be read as one.** It was performed by the same kind of process that produced the proposals under question. It can therefore test the grades against a written rule, but it cannot establish that the rule was applied by a mind uninfluenced by the earlier suggestion.

An independent second assessment by a qualified person remains outstanding, and is named here as an open item rather than quietly omitted. What can be asserted is narrower than the earlier formulation of this note. The grading rule, its application across the whole body of sources, 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.