




INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

FHIR to OMOP Mapping: Loss-Aware Crosswalk and Validation

Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-10-13 · fhir to omop mapping · fhir to omop crosswalk · omop cdm · hl7 fhir · mapping validation · terminology mapping · real world evidence · health data engineering · data quality

Original article: <https://intuitionlabs.ai/articles/fhir-to-omop-mapping-validation>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Key Changes: From Conversion Rules to an Analysis Contract
 4. Resource-to-Table Crosswalk
 5. Terminology, Values, and Vocabulary Permissions
 6. Status, Time, and Provenance: Maintain a Loss Register
 7. Implementation Considerations and Process Changes
 8. Data Analysis and Evidence
 9. Implications and Future Directions
 10. Conclusion
-

Executive Summary

FHIR to OMOP mapping is a controlled transformation from exchange-oriented clinical resources into analysis-oriented tables and standardized concepts. **Fast Healthcare Interoperability Resources (FHIR)** and the **Observational Medical Outcomes Partnership Common Data Model (OMOP CDM)** serve different purposes: interoperability and exchange versus research analytics. A useful crosswalk therefore specifies which meaning an observational analysis needs, where that meaning lands, and what remains outside the target representation. (hl7vulcan.org) A mapping should be accepted for a defined cohort, exposure, outcome, and time window, with a documented loss register, rather than certified as generically lossless. European Medicines Agency (**EMA**) guidance explicitly recognizes that transformation to a common data model can introduce information loss. (www.ema.europa.eu)

The version baseline matters. Reviewed on **October 13, 2026**, the permanent HL7 mapping publication identifies **1.0.0, INFORMATIVE 1**, while the continuous build identifies **2.0.0-ballot** and states that it is not an authorized publication. (hl7.org) (build.fhir.org) The build describes output conforming to **CDM 5.4**, although OHDSI now identifies **5.5** as current. (build.fhir.org) (ohdsi.github.io) These facts require a version manifest, not an assumption that the newest source page, schema, and vocabulary form one tested combination. The manifest should include the declared source profiles and actual FHIR resource release.

The strongest quantitative example reviewed here reports **74% forward mapping coverage** and approximately **23% reverse coverage**. Those are implementation-specific mapping measures, not clinical accuracy, production yield, or evidence that arbitrary records can be reconstructed. (pmc.ncbi.nlm.nih.gov) This report supplies a resource-to-field crosswalk, accountable owners, reconciliation formulas, and an acceptance matrix. The proposed controls are editorial implementation choices; their thresholds must come from the study's needs and source expectations.

The practical decision is whether critical information survives with its clinical meaning. Identity linkage, encounter context, medication evidence type, component observations, units, event status, and temporal precision require separate tests. Unknown modifier extensions need rejection or quarantine because they can change meaning. (build.fhir.org) Vocabulary permissions also require an independent review: inclusion in a

terminology distribution does not automatically authorize all third-party content. (loinc.org) A release is fit for its proposed use only when required mappings are verified, unresolved ambiguities have explicit dispositions, and the study owner accepts the remaining loss.

Introduction and Background

Real-world evidence (**RWE**) platforms often need both clinical exchange and reproducible observational analysis. FHIR supplies resource-based clinical information; OMOP supplies a common analytical representation. Vulcan describes its FHIR-to-OMOP project as a collaboration between HL7 and Observational Health Data Sciences and Informatics (**OHDSI**). (hl7vulcan.org) The distinction between their roles explains why a successful exchange interface does not, by itself, establish that the resulting research database is suitable for a particular study. (hl7vulcan.org)

Mapping, in this report, includes structural placement, terminology translation, value normalization, event selection, and preservation of context. **Extract, transform, load (ETL)** includes those decisions plus ingestion, identity resolution, incremental processing, and loading. The HL7 guide supplies logical models and StructureMaps for transformations, offering a shared technical baseline. (build.fhir.org) This report adds an explicit analysis-purpose acceptance method around that baseline. It does not claim that every local extension or clinical workflow already has an agreed community mapping.

For pharmaceutical teams, the useful question is whether transformations preserve the evidence needed to identify eligible people, distinguish exposure from ordering, and locate outcomes in time. The US Food and Drug Administration (**FDA**) discusses data reliability through accuracy, completeness, and traceability in its electronic-health-record and claims guidance. (www.fda.gov) EMA similarly frames dataset quality assessment as guidance and recommendations. (www.ema.europa.eu) Neither observation establishes that a particular converter or draft mapping artifact is regulator-endorsed.

IntuitionLabs' relevant perspective is that of an adjacent data-engineering and implementation consultancy. Its service description covers pipelines, integration, warehousing, and business intelligence. (intuitionlabs.ai) Its stated information-foundation approach includes identity, permissions, evaluation, and accountable operation. (intuitionlabs.ai) Here, that perspective informs the proposed ownership and evidence controls. It supplies no proprietary conversion benchmark, private deployment result, or competing software product. Readers should use the crosswalk to define a reviewable transformation contract and then test that contract against their own source profiles and analysis requirements.

Key Changes: From Conversion Rules to an Analysis Contract

Pin the publication, resource release, and target schema

The most consequential change in implementation practice is to treat the mapping baseline as a versioned dependency. The published informative artifact and the changing ballot build must be distinguished. (hl7.org) (build.fhir.org) The reviewed guide package identifies a FHIR release, but that does not establish the release of incoming clinical resources. Official US Core ecosystem documentation describes US Core as built on **FHIR**

R4, with versions corresponding to particular US Core Data for Interoperability (**USCDI**) releases.

(isp.healthit.gov) (isp.healthit.gov) A pipeline cannot infer source field names or validation rules solely from the guide's publication package.

The target deserves the same treatment. **CDM 5.5** adds **seven optional fields** and **three vocabulary-metadata tables**, with no fields or tables removed or renamed in the documented release. (github.com) (github.com) (github.com) Compatibility still needs to be demonstrated for the selected loader, mappings, quality checks, and study code. The guide's declared **5.4** output is a separate version statement. (build.fhir.org)

Define the minimum retained meaning

Before implementation, the proposed transformation contract should contain:

- **Analysis:** cohort definition, exposure interpretation, outcome definitions, and exclusions.
- **Source:** FHIR release, profile canonical URLs and versions, extension definitions, and reference-resolution rules.
- **Target:** CDM schema, vocabulary distribution, terminology services, and mapping package identifiers.
- **Identity:** person-linkage scope, encounter identifiers, duplicate resolution, and update behavior.
- **Time:** event-date hierarchy, precision requirements, timezone policy, and permitted derivations.
- **Loss:** retained fields, sidecar fields, excluded records, and unresolved meanings.
- **Acceptance:** study-specific tests, named approvers, and conditions requiring a blocked release.

These are proposed implementation controls. The reviewed continuous build starts with International Patient Access (**IPA**) profiles plus Encounter and Procedure profiles from **US Core STU7**, and says implementers must declare accepted inputs. (build.fhir.org) A critical consequence is that an unrecognized source profile cannot silently inherit a mapping validated for a different profile.

Identity provides a concrete example. The published Encounter map recommends independent integer keys and an external linkage table connecting them to original FHIR identifiers. (hl7.org) The analytical identifier is therefore not a replacement for the original source identity. The OMOP specification also requires an available birth year for PERSON. (ohdsi.github.io) The contract must state the disposition of a source patient without that information, rather than invent a birth year to satisfy a load constraint.

Resource-to-Table Crosswalk

The following crosswalk is an **implementation design reference**, not a claim that every listed resource has a complete executable mapping in the current guide. The guide demonstrates component-observation expansion. (build.fhir.org) Some routes below require local rules, especially medication evidence and document processing.

Table 1 summarizes candidate destinations, vocabulary decisions, expected expansion, reversibility, checks, and ownership. Cardinality describes the proposed transformation's output relationship, not a universal count for every FHIR resource. "Conditional" reversibility requires original source data and lineage. The abbreviated SQL predicates are review checks to adapt to the deployed schema and staging tables.

Resource and target fields	Transform and vocabulary	Cardinality and loss class	Reversibility	Validation query or assertion	Accountable owner
Patient → PERSON.person_id, year_of_birth	Resolve people before allocating stable keys; keep identifiers externally. Patient-to-Person is documented. (build.fhir.org)	Many source identities may resolve to a person; identity ambiguity is conditional loss.	Conditional, with linkage and original demographics.	SELECT person_id, COUNT(*) FROM person GROUP BY person_id HAVING COUNT(*)>1 plus independent linkage review.	Identity lead
Encounter → VISIT_OCCURRENCE.visit_occurrence_id, person_id, dates; VISIT_DETAIL when justified	Define encounter grouping and preserve the source-to-target key relationship externally. (hl7.org)	Source encounters may be split or grouped by declared local rules; encounter-context loss.	Conditional; grouping alone does not reconstruct source encounters.	Anti-join visits against PERSON; compare expected encounter dispositions with loaded visits.	Clinical ETL lead
Condition → CONDITION_OCCURRENCE.condition_concept_id, start date, source fields	Resolve vocabulary domain and distinguish clinical onset from documentation time. (build.fhir.org) Rule-out diagnoses should not populate this table. (ohdsi.github.io)	Exclusion or domain-dependent output; status and onset ambiguity.	Conditional; original verification and date evidence required.	WHERE condition_end_date < condition_start_date; audit the fallback-date flag.	Clinical information
MedicationStatement and other medication evidence → DRUG_EXPOSURE.drug_concept_id, dates, drug_type_concept_id	Separate reported use, prescriptions, dispensing, and administrations. The reviewed executable map excludes MedicationRequest. (build.fhir.org)	Potential exclusion or vocabulary expansion; evidence-type loss if merged.	Conditional; ordering cannot reconstruct administration.	Group staging evidence by resource type and disposition; reconcile resulting exposures separately.	Pharmacoepidemiologist
Procedure → PROCEDURE_OCCURRENCE.procedure_concept_id or domain-appropriate MEASUREMENT/OBSERVATION	Route by the target standard concept's domain. (build.fhir.org)	Domain-dependent output; planned or absent events require explicit handling.	Conditional; clinical event state must survive externally if omitted.	Compare resolved concept domain with destination table; audit excluded event states.	Procedure domain owner
Observation → MEASUREMENT.measurement_concept_id, value_as_number, value_as_concept_id, unit fields, ranges	Preserve code, value semantics, components, and units; review the target standard concept domain. (ohdsi.github.io)	Component observations can expand into multiple measurements. (build.fhir.org) Numeric and context loss remain possible.	Conditional; a component row does not reconstruct its parent resource.	Anti-join component lineage against measurements; compare value and unit pairs before and after transformation.	Laboratory information

Resource and target fields	Transform and vocabulary	Cardinality and loss class	Reversibility	Validation query or assertion	Accountable owner
Other Observation → OBSERVATION.observation_concept_id, value fields, dates	Select the destination from clinical meaning and vocabulary domain rather than resource name alone. Standard concepts have a single domain. (ohdsi.github.io)	Domain-dependent; coded-value and contextual ambiguity.	Conditional; retain original codings and value type.	Audit concept-domain mismatches and contradictory source codings.	Terminology lead
Specimen → SPECIMEN.specimen_concept_id, specimen_date, source fields	Retain sample identity and collection evidence. OMOP describes biological samples; FHIR collection time is optional. (ohdsi.github.io) (hl7.org)	Policy-dependent admission when collection time is missing; sample-context loss.	Conditional; handling and collection details may need a sidecar.	WHERE specimen_date IS NULL; reconcile exclusions against staged samples.	Specimen-data owner
DocumentReference or clinical text → NOTE.note_text, date, type; NOTE_NLP only after extraction	Resolve attachment access and preserve reference/document provenance separately. (hl7.org) (hl7.org)	Attachment- and extraction-dependent; document structure and provenance loss.	Conditional; extracted terms do not recreate the document.	Compare attachment dispositions with loaded note hashes; separately validate extracted assertions.	Document-processing lead

This table makes **destination selection** and **meaning preservation** separate review tasks. A clean destination row can still carry the wrong evidence type or date. FHIR distinguishes administration from ordering and dispensing; those distinctions must inform local exposure rules. (hl7.org) Likewise, a lab result's identity includes more than a display name: LOINC distinguishes properties such as sample, scale, and method. (loinc.org) The table's owners should approve both the clinical interpretation and the corresponding validation evidence.

Terminology, Values, and Vocabulary Permissions

Keep source codes and analytical concepts distinct

Standard-concept coverage measures successful translation into the selected OMOP analytical vocabulary. It is different from preserving the source's original coding. OMOP source-concept fields describe the terminology and concept used in the original data; standard concepts carry domain assignments that inform target placement. (ohdsi.github.io) (ohdsi.github.io) The proposed mapping record should retain original coding system, code, display, version, all competing codings, chosen concept, relationship used, and the reason for any rejection.

A single CodeableConcept can contain alternative representations. The pipeline should not count every coding as an independent clinical event. As an implementation choice, it should select a documented preferred representation and record unresolved disagreement. Vocabulary release changes can alter mappings and

standard-concept properties, so the release belongs in lineage rather than being treated as a background service detail. ([build.fhir.org](#)) Review should focus on study-relevant concepts as well as aggregate coverage. ([ohdsi.github.io](#))

Normalize values with their clinical identity

Logical Observation Identifiers Names and Codes (**LOINC**) expresses observation names through component, property, time, sample or system, scale, and method where relevant. ([loinc.org](#)) The method axis helps distinguish measurements that other axes do not distinguish sufficiently. ([loinc.org](#)) Consequently, matching test display text alone is an inadequate proposed normalization rule.

The Unified Code for Units of Measure (**UCUM**) supports unambiguous exchange of quantities and units. ([ucum.org](#)) Dimensional compatibility does not establish that unit expressions are equal, and arbitrary units can depend on the measurement procedure. ([ucum.org](#)) ([ucum.org](#)) The proposed value rules should therefore include:

- **Quantity:** transform the number and unit together, keeping the original pair.
- **Comparator:** preserve less-than or greater-than meaning instead of turning a bound into an exact value.
- **Reference interval:** apply any verified conversion to the interval as well as the result.
- **Assay context:** block unsupported conversions when comparability depends on method.
- **Coded result:** retain the chosen coded value and the original alternatives.
- **Missing result:** distinguish absent information from an explicit clinical negative.

These are implementation recommendations, not a universal conversion catalog. Unit preservation should be verified with the selected assay and transformation rules; the implementation should not assume that similarly named tests are interchangeable.

Review vocabulary permissions separately

LOINC permits commercial and noncommercial use under its terms, while inclusion of third-party material does not automatically grant permission for that material. ([loinc.org](#)) ([loinc.org](#)) UCUM likewise provides a conditional no-charge license. ([ucum.org](#)) SNOMED International describes no-charge licensed use in Member territories and possible charges in non-Member territories; its vendor guidance requires Affiliate or sublicense status and notes additional Member requirements. ([www.snomed.org](#)) ([www.snomed.org](#)) ([docs.snomed.org](#)) ([docs.snomed.org](#))

The US National Library of Medicine (**NLM**) does not charge for RxNorm licensing, but source-provider content may require additional permissions. ([www.nlm.nih.gov](#)) ([www.nlm.nih.gov](#)) Unified Medical Language System (**UMLS**) terms also recognize additional source restrictions. ([www.nlm.nih.gov](#)) The proposed control is a vocabulary-permission inventory covering deployment territories, redistribution, and included source content. Mapping a code into OMOP does not, by itself, settle those underlying permissions.

Status, Time, and Provenance: Maintain a Loss Register

Treat clinical state as part of the event

The loss register should treat recognized modifier extensions as potentially consequential to event meaning. The guide says data containing unrecognized modifier extensions must not be processed. ([build.fhir.org](#)) The proposed loss register should identify whether each recognized modifier changes event selection, value interpretation, or destination placement. Non-modifier extensions should also have an explicit disposition when the analysis depends on their content.

Status policies need resource-specific review. A procedure marked not-done must not become a procedure occurrence under the reviewed guidance. ([build.fhir.org](#)) Medication-state filters therefore need a pinned acceptance matrix and resource-specific test cases. A prescribing timestamp describes a prescriptive event rather than actual medication initiation. ([build.fhir.org](#)) Implementers should resolve the build's inconsistent medication-status instructions against a pinned rule set and explicit test cases before accepting that branch.

The loss register proposed here uses these classes:

- **Preserved:** clinical meaning survives in declared target fields and verified mappings.
- **Normalized:** meaning survives through a documented code, unit, or value transformation.
- **Externally retained:** meaning needs source storage or a linked sidecar for reconstruction.
- **Derived:** a target value is inferred under a named, approved rule.
- **Ambiguous:** competing interpretations remain unresolved and affect acceptance.
- **Excluded:** a resource or element is intentionally omitted with an accountable disposition.

For every entry, record the source path, target path, rule version, affected analysis variable, reversibility, owner, and acceptance decision. These classes are editorial controls rather than official OMOP categories.

Preserve the distinction between event time and available precision

The guide identifies that OMOP lacks a standard mechanism for original temporal-precision metadata. ([build.fhir.org](#)) The target specification recommends leaving missing datetimes null. ([ohdsi.github.io](#)) Consequently, a year-only birth date or date-only event must not silently acquire the same evidentiary precision as a timestamp. A timezone-normalization policy should retain the original offset and lexical value outside the core representation when needed for reconstruction.

Repeated observations need equally explicit rules: distinguish separate clinical events from successive versions of the same resource. The proposed event identity combines source namespace, resource identity, version or update evidence, and component identity. World Wide Web Consortium (**W3C**) PROV represents revisions as new entities related to earlier revisions, providing a useful lineage pattern. ([www.w3.org](#)) Reprocessing, correction, and replacement should have declared effects on active analytical rows.

Preserve production history, not just resource identifiers

PROV describes provenance through entities, activities, and people, including transformation activities that create entities. ([www.w3.org](#)) ([www.w3.org](#)) USCDI provenance separately describes who created data and when. ([isp.healthit.gov](#)) FHIR DocumentReference also distinguishes provenance of the reference from

provenance of the referenced document. (hl7.org) A proposed sidecar should therefore retain authorship, source version, transformation run, rule and vocabulary versions, disposition, and target keys. That sidecar improves traceability; it does not make a transformed analytical record inherently reversible.

Implementation Considerations and Process Changes

The pipeline should produce evidence at **structural**, **semantic**, and **analysis** levels. EHDEN describes the Data Quality Dashboard (**DQD**) as evaluating and communicating OMOP instance quality, organized around conformance, completeness, and plausibility. (www.ehden.eu) (www.ehden.eu) These dimensions complement source-level validation and sponsor-specific interpretation checks.

The following proposed process gives each release a reviewable record:

- **Ingest:** retain source Bundles or equivalent immutable snapshots and their manifests.
- **Validate source:** test declared release and profile constraints, extensions, and reference resolution.
- **Resolve identity:** separate person linkage from target-key allocation and keep linkage evidence.
- **Resolve terminology:** record accepted, unmapped, ambiguous, and invalid mappings by study variable.
- **Transform:** apply versioned status, date, domain, component, and value rules.
- **Reconcile:** assign each source resource a disposition and an expected target-row expansion.
- **Validate target:** run the matching CDM structural and terminology checks.
- **Validate analysis:** compare cohort membership, exposure interpretation, and outcome timing with source evidence.

The guide allows explicit rejection or quarantine for profile-invalid inputs. (build.fhir.org) In related transformation research, a pipeline reads destination capabilities and validates generated records before writing them. (www.jmir.org) (www.jmir.org) Those practices support a proposed boundary between accepted inputs and clinically interpretable outputs, although that study is not a benchmark for this FHIR-to-OMOP pipeline.

DQD's current index describes **over 20 parameterized check types** that generate individual checks, rather than a fixed universal checklist. (ohdsi.github.io) Run the configured version against the actual target CDM version. (ohdsi.github.io) Domain validation and missing-concept validation answer different questions: a domain check can exclude concept zero from its numerator, so passing it does not demonstrate successful terminology mapping. (ohdsi.github.io) Source-linked optional concept fields may still count as incomplete when their source value exists. (ohdsi.github.io)

Table 2 proposes an acceptance matrix. Its decisions depend on the named analysis and documented evidence; it supplies no universal numerical tolerance.

Decision	Evidence required	Accountable acceptance
Go	Critical variables retain defined meaning; dispositions reconcile; study-level comparisons support the intended analysis.	Study lead and data owner approve the release.

Decision	Evidence required	Accountable acceptance
Go with mitigation	A limitation is bounded, externally retained, or handled through an approved restriction or sensitivity analysis.	Clinical owner approves the mitigation and affected analyses.
Hold	A critical terminology mapping, date rule, or identity ambiguity remains unresolved.	Domain owner supplies evidence before release.
Block for this purpose	Unknown meaning or unavailable essential evidence prevents interpretation of a required study variable.	Study lead records the reason and any narrower permitted use.

The matrix separates loading success from study readiness. OHDSI recommends study-specific quality assessments and reviewing concepts most relevant to the study. (ohdsi.github.io) (ohdsi.github.io) EMA notes that plausibility checks cannot fully replace validation against primary source records. (www.ema.europa.eu) A release should therefore include source review or adjudication appropriate to the critical variables, rather than equating a clean dashboard with clinical fidelity.

Data Analysis and Evidence

Interpret published coverage with its denominator

The reviewed TermX paper reports **74% FHIR-to-OMOP coverage** and approximately **23% reverse coverage**, based on a mapping implementation and synthetic examples. (pmc.ncbi.nlm.nih.gov) (pmc.ncbi.nlm.nih.gov) These figures concern selected mapped structures, not a representative patient population. The authors acknowledge incomplete concept-based semantic testing. (pmc.ncbi.nlm.nih.gov)

These results support a cautious conclusion: forward coverage cannot be interpreted as reverse reconstruction or clinical correctness. They do not supply a universal success threshold. A separate study using data from **two institutions** evaluates legacy-data-to-FHIR transformation; it supports provenance and validation practices, but it is not a FHIR-to-OMOP performance estimate. (www.jmir.org) (www.jmir.org)

Use explicit denominators for local measurements

Table 3 defines proposed metrics. The eligible population and required fields must be declared before measurement. Every percentage is calculated within that declaration, with exclusions reported alongside it.

Metric	Calculation	Interpretation and owner
Mapped-field coverage	Required source-field instances with approved target or sidecar representation ÷ eligible required source-field instances × 100	Excludes optional fields only by prior declaration. Clinical mapping owner reviews critical-field failures.
Unmapped-extension count	Count extension instances with no approved interpretation or disposition; also report distinct canonical URLs	Separate unknown modifiers from non-modifier extensions. Profile owner approves dispositions.
Standard-concept coverage	Eligible coded events assigned valid standard concepts in the expected domain ÷ eligible coded events × 100	Publish domain, source-system, and study-variable strata. Terminology owner checks unmapped and ambiguous events.

Metric	Calculation	Interpretation and owner
Row-count reconciliation	Loaded rows minus the sum of expected output rows across admitted source resources	Expected expansion includes components and declared vocabulary routing. ETL owner investigates each unexplained difference.
Temporal-precision downgrade	Retained events whose target representation has lower time precision $\div$ eligible retained events $\times$ 100	Record original precision and affected study windows. Study lead assesses consequence.
Round-trip non-reversibility	Audited source elements that cannot be reconstructed equivalently $\div$ audited source elements $\times$ 100	Distinguish reconstruction from core tables alone versus core plus retained source. Lineage owner defines equivalence.

The metrics are proposed reporting definitions, not a published standard. DQD's completeness checks count missing field values, and its guidance recommends thresholds customized to source expectations.

(ohdsi.github.io) (ohdsi.github.io) Standard-concept reporting should likewise distinguish genuine missing units from failed mappings: unitless measurements should have null unit concepts rather than concept zero in the reviewed check guidance. (ohdsi.github.io)

Synthetic pipeline check (Hypothetical Example)

An illustrative in-memory test could use a synthetic collection Bundle containing a Patient, an Encounter, a component blood-pressure Observation, and an Observation with an unknown modifier. The proposed setup would use a constrained OMOP-shaped SQLite schema subset to enforce person and visit references. Expected behavior to test would include creation of a person, a visit, and separate systolic and diastolic measurements, quarantine of the unknown modifier, and retention of component lineage. A replay assertion should require that reprocessing the admitted observation adds no rows. This example describes a proposed test design and expected behavior, not reported execution results, production measurements, or full FHIR-profile/CDM validation.

The expansion rule follows the guide's documented component pattern. (build.fhir.org) The quarantine rule follows its unknown-modifier guidance. (build.fhir.org) The proposed checks should cover foreign-key linkage, component expansion, replay behavior, and exclusion disposition. A deployment should extend these tests to its actual schema, profiles, source codings, status rules, and date derivations. OHDSI recommends synthetic ETL unit tests alongside wider quality assessment. (ohdsi.github.io)

Implications and Future Directions

The chief implication for **RWE platform** owners is that **fitness for purpose belongs to a release and an analysis**, not merely to a converter. FDA's EHR and claims guidance is dated **July 2024** and recommends model-conformance checks and documented transformation practices that preserve meaning and context. (www.fda.gov) (www.fda.gov) (www.fda.gov) EMA published the final RWD chapter of its medicines-regulation data-quality framework on **March 27, 2026**, and recognizes possible loss through common-data-model transformation. (www.ema.europa.eu) (www.ema.europa.eu) These sources justify attention to transformation evidence; they do not establish a universal FHIR-to-OMOP mapping mandate.

Future implementation work should distinguish community-guide progress from site-specific analytic validation. A version upgrade should trigger a targeted impact review of changed concepts, resource paths, status logic, time derivations, and target schema. The proposed release evidence should include:

- **Manifest:** immutable mapping identifiers, source-profile packages, target schema, and vocabulary release.
- **Loss register:** critical fields, disposition, reversibility, affected analyses, and accountable owner.
- **Quality results:** structural failures, terminology gaps, reconciliation differences, and study-specific comparisons.
- **Change assessment:** differences against the previous approved release and required reruns.
- **Acceptance record:** approved purposes, mitigations, restrictions, and unresolved items.

Provenance formats such as PROV offer a standards-based way to represent transformation history and revisions. (www.w3.org) (www.w3.org) USCDI authorship metadata can contribute source context, but target row identifiers alone do not capture who created or revised the information. (isp.healthit.gov) The proposed evidence package should preserve both source authorship and pipeline activity.

An OMOP analytical dataset also does not replace every other research representation. Clinical Data Interchange Standards Consortium (**CDISC**) limits its cited observational-study implementation guidance to Study Data Tabulation Model (**SDTM**) scope and publishes example LOINC-to-laboratory mappings for that representation. (www.cdisc.org) (www.cdisc.org) Those are separate transformation contracts. IntuitionLabs' stated focus on accountable information operation is consistent with the ownership model proposed here, but this report claims no private validation evidence or proprietary mapping results. (intuitionlabs.ai)

Conclusion

FHIR-to-OMOP mapping should be approved as an **analysis-specific transformation with declared loss**. The distinction between exchange and analytics explains why an apparently complete set of target tables can still lack context needed for an observational study. (hl7vulcan.org) The useful deliverable is a pinned mapping contract, a field-level crosswalk, a loss register, and evidence that important clinical meaning survives in the selected target representation.

The acceptance decision should address four questions:

- **Identity:** can each analytical person and event be linked to the correct source evidence?
- **Meaning:** do codes, values, units, statuses, and medication evidence represent the intended clinical facts?
- **Time:** do event dates and retained precision support the cohort and outcome windows?
- **Evidence:** can the owners explain exclusions, transformations, ambiguities, and limitations for this release?

Acceptance should include a source-profile declaration; external identifier linkage is a documented part of the mapping approach. (hl7.org) For interpretation, general quality checks should be supplemented by study-specific assessment. (ohdsi.github.io) Information loss and reverse reconstruction need independent measurement, because a forward coverage result cannot settle either question. (www.ema.europa.eu)

The proposed decision framework allows acceptance, mitigation, a hold, or blocking for a particular purpose. It deliberately leaves thresholds to the source expectations and study requirements. Vocabulary permissions remain a separate dependency, including additional restrictions associated with some source content. (www.nlm.nih.gov) The resulting release record should make every consequential choice inspectable: what was preserved, normalized, derived, retained externally, excluded, or left unresolved, and who accepted its effect on the analysis.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://intuitionlabs.ai/articles/introduction-to-fhir-interoperability>
2. <https://intuitionlabs.ai/articles/omop-common-data-model-ohdsi-real-world-data>
3. <https://hl7vulcan.org/projects/fhir-to-omop/>
4. https://www.ema.europa.eu/en/documents/other/data-quality-framework-eu-medicines-regulation-application-real-world-data_en.pdf
5. <https://hl7.org/fhir/uv/omop/INFORMATIVE1/en/StructureMap-EncounterVisitMap.html>
6. <https://build.fhir.org/ig/HL7/fhir-omop-ig/>
7. <https://build.fhir.org/ig/HL7/fhir-omop-ig/en/F2OGeneralIssues.html>
8. <https://ohdsi.github.io/CommonDataModel/>
9. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12935678/>
10. <https://build.fhir.org/ig/HL7/fhir-omop-ig/en/FHIRModifierExtensions.html>
11. <https://loinc.org/kb/license>
12. https://build.fhir.org/ig/HL7/fhir-omop-ig/en/technical_artifacts.html
13. <https://www.fda.gov/media/152503/download?attachment=>
14. <https://www.ema.europa.eu/en/about-us/how-we-work/data-regulation-big-data-other-sources/data-quality-framework-medicines-regulation>
15. <https://intuitionlabs.ai/>
16. <https://isp.healthit.gov/fhir-ecosystem-0>
17. <https://github.com/OHDSI/CommonDataModel/releases>
18. <https://build.fhir.org/ig/HL7/fhir-omop-ig/en/SourceData.html>
19. <https://ohdsi.github.io/CommonDataModel/cdm54.html>
20. <https://build.fhir.org/ig/HL7/fhir-omop-ig/en/StructureMap-PersonMap.html>
21. <https://build.fhir.org/ig/HL7/fhir-omop-ig/en/StructureMap-ConditionMap.html>
22. <https://build.fhir.org/ig/HL7/fhir-omop-ig/branches/main/en/StructureMap-MedicationMap.html>
23. <https://build.fhir.org/ig/HL7/fhir-omop-ig/branches/main/en/StructureMap-ProcedureMap.html>
24. <https://ohdsi.github.io/CommonDataModel/dataModelConventions.html>
25. <https://hl7.org/fhir/R4/specimen.html>
26. <https://hl7.org/fhir/R4/documentreference.html>

27. <https://hl7.org/fhir/R4/medicationadministration.html>
28. <https://loinc.org/kb/faq/structure>
29. <https://ohdsi.github.io/TheBookOfOhdsi/DataQuality.html>
30. <https://ucum.org/ucum>
31. <https://ucum.org/license>
32. <https://www.snomed.org/get-snomed>
33. <https://docs.snomed.org/snomed-ct-practical-guides/vendor-introduction-to-snomed-ct/7-licensing>
34. <https://www.nlm.nih.gov/research/umls/rxnorm/overview.html>
35. https://www.nlm.nih.gov/research/umls/knowledge_sources/metathesaurus/release/license_agreement.html
36. <https://www.w3.org/TR/prov-primer/>
37. <https://www.w3.org/TR/prov-overview/>
38. <https://isp.healthit.gov/uscdi-data-class/provenance?page=1>
39. https://www.ehden.eu/wp-content/uploads/2020/03/EHDEN_Report_Year1_External.pdf
40. <https://www.jmir.org/2023/1/e42822/>
41. <https://ohdsi.github.io/DataQualityDashboard/articles/checkIndex.html>
42. <https://ohdsi.github.io/DataQualityDashboard/articles/checks/cdmField.html>
43. <https://ohdsi.github.io/DataQualityDashboard/articles/checks/fkDomain.html>
44. <https://ohdsi.github.io/DataQualityDashboard/articles/checks/standardConceptRecordCompleteness.html>
45. <https://ohdsi.github.io/DataQualityDashboard/articles/checks/measureValueCompleteness.html>
46. <https://intuitionlabs.ai/articles/real-world-evidence-platforms-fda-ema>
47. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/real-world-data-assessing-electronic-health-records-and-medical-claims-data-support-regulatory>
48. <https://www.cdisc.org/standards/real-world-data>

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore [AI enablement](#), [pharma and biotech software development](#), [data engineering and BI](#), and [Veeva services](#). [Contact IntuitionLabs](#) to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: <https://intuitionlabs.ai>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.