




INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

FDA PQ/CMC FHIR: Module 3 Architecture and ICH M16 Guide

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 · 2026-10-07 · fda pq/cmc fhir · pq/cmc implementation guide · structured cmc data · ich m16 · ectd module 3 · idmp · fhir r5 · quality data architecture · regulatory data governance

Original article: <https://intuitionlabs.ai/articles/fda-pq-cmc-fhir-data-architecture>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Key Changes
 4. Implementation Considerations and Process Changes
 5. Canonical Quality Data and Provenance
 6. PQ/CMC and IDMP: Semantic Overlap and Governance
 7. Validation and Controlled Document/Payload Generation
 8. Data Analysis and Evidence
 9. Implications and Future Directions
 10. Conclusion
-

Executive Summary

FDA PQ/CMC FHIR is a technical approach to expressing selected pharmaceutical quality, chemistry, manufacturing, and controls information as structured, computable data. As of **October 7, 2026**, the published Health Level Seven International (**HL7**) implementation guide identifies itself as **version 2.0.0, Standard for Trial Use 2 (STU2)**, built on **Fast Healthcare Interoperability Resources (FHIR) 5.0.0**. The continuous-build site also displays 2.0.0, but expressly says it is not an authorized publication. Identical version labels therefore do not make those two artifacts interchangeable. (hl7.org) (build.fhir.org)

The regulatory boundary is equally consequential. FDA states that HL7 testing, balloting, and publication do not constitute FDA policy or guidance. (www.fda.gov) Its **June 2026** assessment says the latest PQ/CMC iteration is complete and anticipates **ICH M16 Structured Product Quality Submissions** carrying the objective forward for original applications and postapproval changes. (www.fda.gov) FDA's **September 2026 Joint Action Plan (v2.4)** records that the **December 2025** update removed the **Pharmaceutical Quality/Chemistry, Manufacturing, and Controls Data Standardization Project** from the action plan. (www.fda.gov) This standalone action-plan status is distinct from anticipated ICH M16 work. The published **STU2** guide remains available. (hl7.org) Removal from the action plan alone does not establish withdrawal of the guide. These sources establish technical progress and future direction; they do not establish a current universal requirement to submit Module 3 as FHIR.

For implementation, **Stage 1** covers substance general information, control of materials, specifications, and drug-product description/composition; **Stage 2** extends substance characterization, product batch formula, and product impurities. The following inventory connects those sections to actual profiles and likely source systems. It distinguishes these delivered stages from FDA's broader conceptual phases, which also include batch analysis and stability. (www.fda.gov) (www.fda.gov) **Identification of Medicinal Products (IDMP)** supplies overlapping identity semantics, while the European Medicines Agency's **SPOR** services organize substance, product, organisation, and referential master data. Reuse is appropriate where definitions match, not where labels merely resemble one another. (www.ema.europa.eu) (globalforum.diaglobal.org)

The recommended investment is a **governed canonical quality-data layer** that preserves source identifiers, approved versions, terminology mappings, units, and derivation history, then produces both documents and testable structured payloads. The FAIR principles support persistent identification and provenance; the World Wide Web Consortium's provenance model supplies a useful vocabulary for tracing entities, activities, and responsible agents. Neither is an extra PQ/CMC submission mandate. (www.gofair.foundation) (www.w3.org) Build identity resolution, specification semantics, and review controls now; pilot pinned FHIR packages; defer assumptions about final M16 packaging and acceptance dates. The readiness calculations below measure evidenced concept and mapping coverage, rather than assigning an unsupported industry-wide readiness percentage. The central decision is whether the organization can reproduce an approved quality statement from governed source data and explain every transformation.

Introduction and Background

The **electronic Common Technical Document (eCTD)** provides submission organization and lifecycle context. The International Council for Harmonisation (**ICH**) describes the Common Technical Document (**CTD**) as **five modules**, with the first region specific and the remaining modules intended to be common. (admin.ich.org) PQ/CMC addresses selected quality information currently submitted in **Module 3 and Module 2.3**, the Quality Overall Summary (**QOS**). FDA explicitly excludes comprehensive structuring of all product-quality information from the project's scope. (www.fda.gov) (www.fda.gov)

This distinction separates three decisions: what scientific evidence belongs in a dossier, how a structured model represents selected concepts, and what a regulator formally accepts. A PDF section, a FHIR profile, and an agency acceptance announcement answer different questions. Preparing a structured representation does not by itself justify changing an established submission workflow. FDA says support for accepting a particular format follows formal announcement processes. (www.fda.gov)

The architecture challenge is consequently broader than an export utility. A specification can exist as an approved document, a laboratory configuration, and a regional submission statement, each with a different effective date. An integration design should preserve those distinctions while identifying the common scientific object. DIA's discussion of data-standards implementation highlights older siloed systems and emphasizes that the published PQ/CMC-IDMP overlap diagram is high-level alignment rather than an element-level mapping. (globalforum.diaglobal.org) (globalforum.diaglobal.org)

The adjacent consultancy perspective is relevant here: **IntuitionLabs** describes connecting authoritative enterprise sources through identity, permissions, retrieval, citations, and evaluation. (intuitionlabs.ai) Applied to quality submissions, that perspective favors governed source reuse and accountable transformations. It does not make a consultancy a competing data standard or a submission authority. This report therefore concentrates on the architecture and decision criteria that regulatory CMC leaders, quality-data stewards, master-data owners, and integration engineers can inspect and test.

Key Changes

Separate conceptual phases from delivered FHIR stages

FDA's conceptual **Phase 1** includes drug substance, drug product, quality specification, batch formula, batch analysis, and stability; conceptual **Phase 2** includes drug-product and drug-substance manufacturing.

(www.fda.gov) (www.fda.gov) The FHIR stages implement subsets incrementally. Treating every Phase 1 concept as a delivered Stage 1 profile would overstate current coverage. (www.fda.gov)

Table 1 maps the delivered section families to representative profile/resource pairs. Profile names and section coverage come from the published guide; source-system assignments are this report's proposed architecture, subject to local ownership decisions.

Stage and CTD section	Structured concept and representative profile	FHIR resource	Proposed authoritative source
Stage 1: 3.2.S.1	Substance nomenclature/structure: DrugSubstanceNomenclatureStructure, PolymorphicForm (hl7.org)	SubstanceDefinition	Substance master; approved development records
Stage 1: 3.2.S.2.3	Materials: ExcipientRaw, SubstanceDefinitionHandle; QualitySpecification (hl7.org)	SubstanceDefinition; PlanDefinition	Material master; supplier records; specification management
Stage 1: 3.2.S.4.1, 3.2.P.4, 3.2.P.5.1	Specifications: QualitySpecification, including tests, stages, acceptance criteria (hl7.org)	PlanDefinition	Approved specification repository; laboratory method registry
Stage 1: 3.2.P.1	Description/composition: DrugProductDescription, ContainerClosure, FinishedProduct, DrugProductComponent, ComponentSubstance (hl7.org)	MedicinalProductDefinition; PackagedProductDefinition; ManufacturedItemDefinition; Ingredient; SubstanceDefinition	Product master; formulation records; packaging specifications
Stage 2: 3.2.S.3	Characterization: DrugSubstanceCharacterisation, ImpuritySubstance (hl7.org)	SubstanceDefinition	Analytical development; substance and impurity registry
Stage 2: 3.2.P.3.2	Batch formula: BatchFormulaMedicinalProduct, BatchFormula, DrugProductIngredient (hl7.org)	MedicinalProductDefinition; ManufacturedItemDefinition; Ingredient	Approved master recipe; manufacturing systems
Stage 2: 3.2.P.5.5	Product impurities: DrugProductWithImpurities, ImpuritySubstance (hl7.org)	MedicinalProductDefinition; SubstanceDefinition	Impurity characterization records; product-quality assessment

The table should become a requirements inventory, not a promise that every local object has a direct resource equivalent. For example, the substance-general-information bundle explicitly excludes **3.2.S.1.3**, despite the broader section label. ([hl7.org](#)) Likewise, the published home page limits scope to **solid oral dosage forms**. A program should document supported modality and section scope before extrapolating a pilot to an entire portfolio. ([hl7.org](#))

Distinguish exchange resources from enterprise objects

The guide uses **document-type bundles**, with identifiers for bundle entries and referenced resources packaged together; its general instructions explain that the receiving endpoint is not a FHIR server. ([hl7.org](#)) FHIR adoption therefore does not imply that a sponsor needs a clinical-style operational FHIR server or a live transaction interface to FDA. The published guide aligns its profiles with **eCTD v4.0 or later** and warns that they may not fit **eCTD v3 headings**. ([hl7.org](#)) Document the pilot's publishing version and compatibility with that boundary. Keep this compatibility check separate from the assessment of formal FDA acceptance.

A recommended implementation keeps enterprise identifiers and approved business versions independent of transport identifiers. The export adapter should express the approved object graph in the required bundle structure. This separation follows the FAIR emphasis on persistent identification and detailed provenance, while leaving packaging replaceable. ([www.gofair.foundation](#)) It also prevents a regulatory submission identifier from becoming the sole identity of a substance or specification throughout the enterprise.

Implementation Considerations and Process Changes

Establish ownership before building mappings

The following source map is a **proposed operating model**, not a statement that a particular application category universally owns each fact. EMA's master-data approach supports entering organisation and referential data once and reusing them across procedures. ([www.ema.europa.eu](#)) Use that principle to resolve authority field by field, especially where manufacturing, laboratory, and regulatory records contain different representations of the same concept.

- **LIMS:** A laboratory information management system should supply approved test results, sample identity, analytical procedure references, and result status where it is authoritative. Test names alone are insufficient: LOINC's observation model distinguishes component, property, timing, system, scale, and sometimes method. That is a semantic design lesson, not a claim that LOINC codes must replace PQ/CMC test terminology. ([loinc.org](#))
- **MES/eBR:** A manufacturing execution system or electronic batch record should supply executed production facts. Keep those separate from the approved master recipe and from the formula selected for submission. FDA describes batch formula as the recipe for making the product. ([www.fda.gov](#))
- **QMS:** A quality management system should supply approval state, effective dates, and change-control references when those records govern release. The design should retain who authorized a transformation, using provenance agents and activities as a useful conceptual model. ([www.w3.org](#))

- **ERP:** An enterprise resource planning system may own material codes, suppliers, and operational sites. Reconcile these identifiers with organisation and referential master data rather than treating every supplier label as a new legal organisation. (www.ema.europa.eu)
- **Specification management:** The approved specification repository should own tests, acceptance rules, applicability, and versions. Retain hierarchy among tests, stages, and acceptance criteria instead of reducing the model to a flat table of test names and limits.
- **RIM:** Regulatory information management should own dossier context, markets, application references, and the approved regulatory representation. EMA's product service identifies medicinal products through regulated information, illustrating why the regulatory identity needs an explicit master-data relationship. (www.ema.europa.eu)
- **eCTD publishing:** The publishing environment should consume approved submission artifacts and preserve lifecycle context. It should not become the only place where scientific facts are maintained. Configure publishing as an artifact consumer, with an explicitly bounded structured-export scope.

Assign a **data owner** accountable for meaning and release, a **steward** responsible for mapping and exception resolution, and a **system custodian** responsible for interfaces and access. These roles are proposed controls. FAIR accessibility can include authentication and authorization; it does not require unrestricted disclosure of proprietary quality data. (www.gofair.foundation)

Build the section-to-element crosswalk

For each document statement, identify the canonical object, its source version, its structured destination, and any supporting attachment. A batch-formula statement must distinguish component amount per batch, overage, and the relevant quality standard: the CTD quality guidance explicitly calls for per-batch amounts including overages. (www.fda.gov) A product-unit composition amount must remain a different field from a manufacturing batch amount.

Store the mapping as a governed record with transformation logic, terminology version, applicability, evidence, and reviewer approval. Use a persistent identifier for the mapping itself, not only the generated file. This is an architectural application of FAIR identification and provenance principles. (www.gofair.foundation) Where a concept has no delivered profile, retain it in the canonical model and document output; mark structured export coverage as a gap rather than inventing a destination.

Canonical Quality Data and Provenance

Model scientific distinctions explicitly

The proposed canonical model should be an enterprise quality model with versioned relationships, not a copy of a single submission schema. ISO's substance standard provides an information model for defining and identifying substances, while NCATS describes **GSRS**, the Global Substance Registration System, as supporting exchange of substance records. Those are useful identity foundations; neither establishes an enterprise's complete quality model. (www.iso.org) (gsrs.ncats.nih.gov)

- **Substance identity:** Preserve substance definition, names, external identifiers, salt or form distinctions, and evidence for equivalence. Do not merge records solely because a display name matches. (www.iso.org)
- **Product composition:** Separate the finished product, its physical components, and ingredient constituents. The IG distinguishes components such as layers or coatings from active and inactive constituents. (hl7.org)
- **Specification version:** Represent the subject, applicability, tests, test stages, acceptance criteria, analytical references, approval state, and effective interval. Preserve their hierarchy before exporting it.
- **Quantitative expression:** Store value, comparator, unit, basis, and denominator explicitly. UCUM, the Unified Code for Units of Measure, distinguishes incompatible case-sensitive and case-insensitive symbol representations; unit normalization must preserve the selected convention. (ucum.org) ISO 11240 also addresses machine-readable quantitative composition and strength. (www.iso.org)
- **Batch and recipe:** Keep planned formula, executed batch, sample, and analytical result as distinct objects linked through stable identifiers. FDA's high-level comparison describes CMC batch needs extending to batch size and manufacturing/testing sites. (www.fda.gov)
- **Reference context:** Preserve reference identity, assigned-value basis, and the applicable measurement chain as separate fields. NIST defines metrological traceability through a calibration chain contributing measurement uncertainty. (www.nist.gov)
- **Scientific evidence:** Link conclusions to the approved records and attachments supporting them. NIST's definition of measurement traceability requires a documented calibration chain contributing measurement uncertainty; a file link alone does not express that scientific chain. (www.nist.gov)

The model should allow **unknown**, **not applicable**, **not yet assessed**, and **confirmed absent** to remain distinct business states. Any export must map them according to the destination profile rather than silently coercing them into empty strings or zero values. An impurity fixture should test this distinction directly: a deliberate absence assertion must remain distinguishable from missing input.

Preserve lineage through generation and review

The World Wide Web Consortium (**W3C**) provenance ontology, **PROV-O**, distinguishes entity, activity, and agent, and describes derivation as transformation of one entity into another. (www.w3.org) (www.w3.org) A practical lineage chain can therefore record an approved source entity, extraction activity, mapping activity, canonical snapshot, generation activity, reviewer, and released artifact. This is a proposed use of provenance concepts, not a recommendation to submit PROV-O to FDA.

For every released artifact, retain the source record identifiers and versions, extraction timestamp, mapping version, terminology snapshot, adapter version, review decision, and artifact checksum. Access to these records should follow the same permission boundaries as the underlying evidence. FAIR explicitly accommodates authorization procedures. (www.gofair.foundation) Reproduction should mean regenerating the same scientific content from the same approved inputs, with any intentional presentation differences explained.

Source lineage and **metrological traceability** must remain separate. The former explains where a data statement came from; the latter connects a measurement to its reference. NIST also cautions that traceability alone does not guarantee fitness for purpose. (www.nist.gov) A complete audit trail cannot establish that an unsuitable analytical procedure supports a specification. Scientific review must remain an independent control.

PQ/CMC and IDMP: Semantic Overlap and Governance

FDA explicitly states that PQ/CMC is **not FDA's implementation of ISO IDMP**. (www.fda.gov) IDMP's medicinal-product standard addresses unique identification and detailed description; the pharmaceutical-product standard also cautions that it is not a scientific classification. (www.iso.org) (www.iso.org) These statements support reuse of identity semantics without equating all product-quality evidence with identification data.

Table 2 distinguishes candidate reuse from mappings requiring further scientific and regulatory interpretation. It is a proposed architecture informed by official scope descriptions, not a certified equivalence crosswalk.

Quality-data area	Identity or terminology overlap	Recommended governance boundary
Substance nomenclature and structure	FDA identifies direct overlap with ISO 11238. (www.fda.gov)	Reuse verified identity mappings; preserve regional identifiers and form distinctions.
Product identity and composition	ISO 11615 addresses medicinal-product description; EMA PMS uses regulated product information. (www.iso.org) (www.ema.europa.eu)	Separate common identity from market-specific regulatory representation and submission context.
Dose form, route, packaging	ISO 11239:2023 supports these concepts and concept versioning; EDQM Standard Terms include containers, closures, devices, and units of presentation. (www.iso.org) (standardterms.edqm.eu)	Govern external terminology versions and regional applicability; do not equate labels without checking definitions.
Quantitative composition and strength	ISO 11240 addresses machine-readable quantitative composition and strength. (www.iso.org)	Preserve amount basis, denominator, units, and transformations.
Specification, test, acceptance criterion	Model tests, analytical references, and acceptance criteria as explicit quality-control rules. Conditions on an optional RDF layer can be expressed using SHACL. (www.w3.org)	Keep quality-control rules as governed quality objects; do not derive limits from product identity alone.
Batch recipe and stability plan	FDA describes batch formula as a recipe and stability information as a structured study plan. (www.fda.gov) (www.fda.gov)	Maintain lifecycle evidence separately; map only demonstrated common elements.

The distinction prevents two common architecture errors: duplicate master-data creation and excessive semantic consolidation. FDA's comparison is a high-level historical analysis, not a certification that every current IDMP implementation lacks a particular quality attribute. DIA likewise warns that the overlap diagram does not represent data-element details. (globalforum.diaglobal.org) Require mapping evidence for each element, including definition, cardinality, context, units, and permitted terminology.

EMA's SPOR model provides a useful organisation of master-data responsibilities across substance, product, organisation, and referential domains. (www.ema.europa.eu) Implement crosswalks as controlled relationships with confidence, evidence, owner, and review status. GSRS record exchange can support substance collaboration, but the fetched evidence does not establish a universal UUID-to-UNII conversion rule; identifier resolution must use documented, applicable services. (gsrs.ncats.nih.gov)

Terminology governance needs its own release process. The ISO dose-form model explicitly includes concept versioning. (www.iso.org) EDQM makes Standard Terms access free with registration, while access terms and downstream reuse rights should be checked separately. (standardterms.edqm.eu) Publicly readable ISO catalog abstracts describe scope; they do not supply the full normative standards or blanket reproduction rights. The practical rule is to inventory each incorporated standard, code system, supplement, and local vocabulary, then record its license, version, owner, and permitted use.

Validation and Controlled Document/Payload Generation

Validate meaning as well as conformance

FHIR validation is necessary but bounded. Official FHIR documentation distinguishes technical checks from the full set of real-world implementation concerns. (hl7.org) The release gate should therefore assess syntax, profile conformance, terminology, referential consistency, scientific semantics, and human review as separate dimensions. A validator pass should never be presented as evidence that FDA accepted a submission.

The following is a proposed test catalog. Each test should have an approved fixture, expected result, retained evidence, and accountable reviewer.

- **Package reproducibility:** Pin the FHIR release, IG package, dependencies, terminology artifacts, and validator build. Retain the package inventory with the generated artifact. The official validator project supplies validation libraries and command-line tooling. (github.com)
- **Profile conformance:** Test required elements, cardinalities, resource types, constraints, and reference targets. The IG download provides a validator pack intended to support resource validation. (hl7.org)
- **Terminology membership:** Test code/system combinations against the applicable binding and version. QualitySpecification requires codes from Test Category Terminology for the relevant element; an approved local display label alone is insufficient. (hl7.org)
- **Supplement loading:** Confirm that the terminology environment includes required supporting artifacts. The Test Category value set expressly requires the relevant code-system supplement. (hl7.org)
- **Numeric semantics:** Test comparators, ranges, precision, unit case, and unit conversions. Preserve the selected UCUM symbol convention throughout conversion and export. (ucum.org)
- **Constraint edge cases:** Include expected failures, such as incompatible simultaneous structures. The specification profile prohibits detail when its target-range extension is present. (hl7.org)
- **Identity resolution:** Test duplicate and ambiguous substances, site identities, and external identifiers. ISO 11238's substance-identification purpose supplies the conceptual boundary, while local equivalence decisions still require evidence. (www.iso.org)

- **Analytical context:** Test sample, method version, result basis, and reference context. LOINC includes a method axis when other axes do not sufficiently distinguish a measurement. (loinc.org) Retain the applicable calibration/reference chain as well. (www.nist.gov)
- **Round-trip preservation:** Generate a payload, parse it back into a comparison model, and compare scientific content, including units, basis, context, and attachments. Preserve derivation links for each transformation. (www.w3.org)
- **Review-state control:** Prevent draft or superseded inputs from entering a released artifact without an explicit authorized decision. Use risk-proportionate validation; ISPE describes GAMP guidance as non-prescriptive rather than a universal method or standard. (ispe.org)

W3C's **Shapes Constraint Language (SHACL)** validates Resource Description Framework (**RDF**) graphs against conditions and can produce a validation report. (www.w3.org) It can support an optional graph-based canonical layer, but it does not replace FHIR validation and does not imply that an RDF layer is required. Use the validation mechanism appropriate to the canonical representation and retain detailed findings rather than only a boolean outcome. (www.w3.org)

Generate documents and payloads from an approved snapshot

Release document and structured outputs from the same approved snapshot, with separate adapter and presentation versions. Compare important human-readable statements directly with the underlying structured facts. The IG's Narrative Transform generates narrative and explicitly does **not** validate the bundle. (hl7.org) Narrative generation and conformance validation must therefore remain separate operations.

Maintain a rights inventory alongside the technical package inventory. The FHIR R5 specification uses **CC0 1.0**, but its license does not grant rights to third-party intellectual property. (hl7.org) NCIt's Terms of Use specify **CC BY 4.0**. (evs.nci.nih.gov) Avoid a blanket claim that every dependency inherits the same permission merely because an IG page is public. Similarly, library support for a serialization does not establish regulatory acceptance of that serialization. The IG's general instructions distinguish XML for downstream submission use from JSON for other implementation uses. (hl7.org)

Data Analysis and Evidence

What the public evidence quantifies

The verified milestones measure standards development, not sponsor readiness or business benefit. FDA's June 2026 assessment records Stage 2 testing at the **September 2024 Connectathon** and completion of the latest PQ/CMC iteration. (www.fda.gov) ICH's **November 2025** Assembly minutes record endorsement of the M16 concept-paper outline in **October 2025**. (admin.ich.org) The published guide and continuous build both currently display **2.0.0**, but only the former identifies itself as the current published version. (hl7.org) (build.fhir.org) A standards-development milestone is not an adoption deadline.

Terminology currency is also measurable. NCI states that the **NCI Thesaurus (NCIt)** is published monthly, so validation environments should record the actual terminology snapshot rather than assuming the latest label is reproducible. (www.cancer.gov) ISO's current catalog identifies **ISO 11239:2023**, not the older edition, for dose

forms, routes, presentation units, and packaging concepts. (www.iso.org) Measurement context also needs preservation: LOINC makes method an optional axis of a fully specified observation name. (loinc.org) These are concrete reasons to capture dependency versions in every pilot.

The reviewed sources do not provide a comparable public dataset of sponsor implementation costs, processing-time reductions, mapping accuracy, or FHIR-based CMC acceptance rates. This report therefore does not convert technical progress into an efficiency percentage. Any local performance claim should identify the source population, measurement interval, baseline workflow, and review outcome.

Calculate readiness from an evidenced inventory

The following measures are **analyst-defined calculations**, not FDA or ICH scoring systems. Count only the concepts and systems within the declared pilot scope. Attach evidence to each completed item so that coverage cannot rise merely by relabeling a field.

- **Concept coverage:** Approved, semantically mapped concepts divided by in-scope concepts.
- **System coverage:** Connected authoritative source systems divided by source systems needed for scope.
- **Mapping coverage:** Approved source-to-canonical transformations divided by required transformations.
- **Term coverage:** Validated controlled-term mappings divided by controlled-term mappings needed.
- **Lineage coverage:** In-scope released statements with complete source/version evidence divided by released statements assessed.
- **Scope stability:** Change a denominator only through an approved scope revision; retain the earlier inventory for comparison.

(Hypothetical Example) Consider an inventory with **120 concepts, 8 systems, 180 transformations**, and **240 controlled-term mappings**. If **90 concepts, 5 systems, 108 transformations**, and **168 term mappings** have approved evidence, coverage is respectively **75%, 62.5%, 60%, and 70%**. These percentages are arithmetic from the stated hypothetical counts, not observed industry figures. The lowest coverage dimension is 60%; averaging the dimensions would obscure a transformation backlog.

For a critical specification, add a release condition: every mandatory input and semantic test must pass, even if portfolio coverage is high. NIST's warning that traceability alone does not establish fitness for purpose applies to interpreting such scores: documentation completeness is not scientific suitability. (www.nist.gov) Pair coverage with exception severity, aging, and unresolved ownership. A management dashboard should make the limiting dependency visible, rather than converting unrelated completion ratios into a single compliance score.

Implications and Future Directions

ICH's current multidisciplinary index names M16 Structured Product Quality Submissions, while FDA's latest assessment anticipates international continuation of the structured-quality objective. (admin.ich.org) (www.fda.gov) This research did not verify a final M16 technical specification or an FDA mandatory-adoption date. Planning should preserve the ability to update mappings and packaging when official requirements become available.

Table 3 sets out a proposed build decision. It separates capabilities with present operational value from technical experiments and decisions dependent on future announcements.

Decision	Scope of investment	Evidence required before expansion
Build now	Govern identity, approved specification semantics, units, terminology versions, source ownership, and lineage. Support current document generation. FAIR provenance and EMA master-data reuse provide relevant design principles. (www.gofair.foundation) (www.ema.europa.eu)	Approved ownership register; reproducible source snapshots; reviewed mappings; document-to-data reconciliation.
Pilot	Export a bounded solid-oral scope against pinned published PQ/CMC artifacts. Exercise profile, terminology, narrative, and round-trip tests. The current IG states its modality limitation. (hl7.org)	Reproducible package inventory; negative-test evidence; scientific review; explicit coverage gaps.
Defer requirement-dependent decisions	Final M16-specific packaging, compulsory portfolio migration, and universal replacement of existing submission workflows. FDA acceptance announcements follow formal processes. (www.fda.gov)	Published requirements, applicable agency acceptance instructions, transition arrangements, and portfolio impact assessment.

The table favors investments that serve both today's quality documentation and later structured submission. It also makes the stop conditions visible: an unresolved identifier, unsupported modality, or ambiguous acceptance criterion should remain a reviewed exception rather than a guessed export. A pilot is valuable when it exposes these dependencies and yields reusable mapping evidence.

The consultancy role is complementary to source applications and regulatory systems. IntuitionLabs describes implementation, extension, integration, and support for Veeva commercial and development applications. (intuitionlabs.ai) In the architecture proposed here, such integration expertise can support the governed data layer without making a consultancy the system of record or a competing standard. Application selection should follow demonstrated source authority, integration behavior, and review needs.

Governance should also preserve explicit rights and terminology-access decisions. EDQM's registered access model and NCI's published terms illustrate why retrieval, reuse, attribution, and redistribution should be assessed separately. (standardterms.edqm.eu) (evs.nci.nih.gov) Maintain a change register for official FDA announcements, published IG artifacts, and ICH materials; assess each change against the canonical model, mappings, validation fixtures, and downstream documents before release.

Conclusion

FDA PQ/CMC FHIR is sufficiently concrete to support a bounded architecture and validation pilot, while the distinction between technical publication and regulatory acceptance remains central. Published and continuous-build artifacts require different treatment, conceptual phases should not be confused with delivered stages, and current modality and section coverage should be recorded explicitly.

The strongest preparation strategy is to govern quality facts at their source: define stable identities, preserve scientific context, approve mappings, version terminology, and retain derivation evidence. IDMP and SPOR can inform common identity and referential governance, but they do not justify collapsing every quality object

into a product-identification record. Their reuse value depends on demonstrated semantic equivalence. (www.ema.europa.eu) (globalforum.diaglobal.org)

Source-system teams should prioritize reproducible approved snapshots and document-to-data reconciliation. Integration teams should keep canonical meaning separate from FHIR packaging and retain evidence for negative tests, round trips, and review decisions. Scientific owners should verify that the result is fit for purpose, alongside the technical team's conformance checks. (www.nist.gov)

The build decision can then be made on observed evidence: build governed reusable quality data, pilot the supported technical scope, and reserve requirement-dependent decisions for official announcements. This approach creates useful capabilities before ICH M16 matures while keeping the eventual submission adapter and transition plan adaptable. The deliverable is an explainable quality statement that can generate an approved document and a tested structured representation from the same governed evidence.

Source links

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

1. <https://hl7.org/fhir/us/pq-cmc-fda/>
2. <https://build.fhir.org/ig/HL7/FHIR-us-pq-cmc-fda/>
3. <https://www.fda.gov/industry/fda-data-standards-advisory-board/pharmaceutical-quality-chemistry-manufacturing-controls-pqcmc>
4. <https://www.fda.gov/media/193422/download>
5. <https://www.fda.gov/media/194922/download?attachment=>
6. <https://www.fda.gov/industry/pharmaceutical-quality-chemistry-manufacturing-controls-pqcmc/pqcmc-and-hl7-fhir>
7. <https://www.fda.gov/industry/pharmaceutical-quality-chemistry-manufacturing-controls-pqcmc/pqcmc-data-standards-scope-development>
8. <https://www.ema.europa.eu/en/human-regulatory-overview/research-development/data-medicines-iso-idmp-standards-overview/substance-product-organisation-referential-spor-master-data>
9. <https://globalforum.diaglobal.org/issue/february-2025/data-standards-fatigue/>
10. <https://www.gofair.foundation/fair-principles>
11. <https://www.w3.org/TR/prov-o/>
12. <https://admin.ich.org/page/ctd>
13. <https://intuitionlabs.ai/>
14. <https://hl7.org/fhir/us/pq-cmc-fda/eCTD32S10.html>
15. <https://hl7.org/fhir/us/pq-cmc-fda/eCTD32S23.html>
16. <https://hl7.org/fhir/us/pq-cmc-fda/eCTDSP4151.html>
17. <https://hl7.org/fhir/us/pq-cmc-fda/eCTD32P10.html>
18. <https://hl7.org/fhir/us/pq-cmc-fda/eCTD32S3.html>
19. <https://hl7.org/fhir/us/pq-cmc-fda/eCTD32P32.html>
20. <https://hl7.org/fhir/us/pq-cmc-fda/eCTD32P55.html>
21. <https://hl7.org/fhir/us/pq-cmc-fda/instruction.html>

22. <https://loinc.org/kb/faq/structure>
23. <https://www.fda.gov/industry/pharmaceutical-quality-chemistry-manufacturing-controls-pqcmc/pqcmc-and-idmp>
24. <https://intuitionlabs.ai/articles/ectd-publishing-process-guide>
25. <https://www.fda.gov/media/71581/download>
26. <https://www.iso.org/standard/69697.html>
27. <https://gsrs.ncats.nih.gov/about>
28. <https://ucum.org/ucum>
29. <https://www.iso.org/standard/55033.html?browse=tc>
30. <https://www.nist.gov/metrology/metrological-traceability>
31. <https://www.iso.org/standard/70150.html>
32. <https://www.iso.org/standard/70044.html>
33. <https://www.iso.org/standard/81133.html>
34. <https://standardterms.edqm.eu/>
35. <https://www.w3.org/TR/shacl/>
36. <https://hl7.org/fhir/R5/validation.html>
37. <https://github.com/hapifhir/org.hl7.fhir.core>
38. <https://hl7.org/fhir/us/pq-cmc-fda/downloads.html>
39. <https://hl7.org/fhir/us/pq-cmc-fda/StructureDefinition-pqcmc-quality-specification-definitions.html>
40. <https://hl7.org/fhir/us/pq-cmc-fda/ValueSet-pqcmc-test-category-terminology.html>
41. <https://ispe.org/topics/gamp>
42. <https://hl7.org/fhir/R5/license.html>
43. https://evs.nci.nih.gov/ftp1/NCI_Thesaurus/ThesaurusTermsofUse.pdf
44. https://admin.ich.org/sites/default/files/2026-01/ICH51_Singapore_Assembly_Minutes_Meeting_Final_2026_1201.pdf
45. <https://www.cancer.gov/about-nci/organization/cbiit/vocabulary>
46. <https://admin.ich.org/page/multidisciplinary-guidelines>

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.