

EMA ePI FHIR v1.0 Implementation Readiness Guide

IntuitionLabs.ai · Adrien Laurent · 2026-09-19 · ema epi fhir v1.0 · electronic product information · eu epi common standard · plm portal · regulatory submissions · fhir validation · product information · centralised procedures · regulatory operations



Executive Summary

The European Medicines Regulatory Network (EMRN) electronic product information (ePI) Implementation Guide **1.0.0** is the first released set of EMRN ePI resource definitions and profiles, dated **17 September 2026** (epi.ema.europa.eu). It is based on **FHIR 5.0.0** (epi.ema.europa.eu). The release makes the technical baseline stable enough for controlled implementation work. It does not mean that routine production submission has begun. As of **19 September 2026**, EMA says its applicant guidance was published before go-live and becomes applicable at go-live (www.ema.europa.eu), while the draft roadmap places phased centrally authorised product go-live from **Q4 2026** (www.ema.europa.eu).

The data model should be implemented as a hierarchy, not as a styled document conversion. A persistent **List** identifies the ePI master record (epi-dev.ema.europa.eu); each SmPC, Annex II, labelling, or package-leaflet document is a **Bundle** of type document (epi.ema.europa.eu); its first entry is a **Composition** (hl7.org); and

images are contained **Binary** resources (epi-dev.ema.europa.eu). CAP and NAP use separate Composition constraints and section-code value sets. The final CAP set has **134 concepts**, compared with **122** in the NAP-MRP-DCP set (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu).

The main architecture decision is whether the regulatory team authors directly in the Product Lifecycle Management (PLM) Portal or imports FHIR XML generated by an internal or supplier platform. Both are documented options (plm-portal.ema.europa.eu). Portal authoring offers the shortest control path, but pasted Word formatting is removed (plm-portal.ema.europa.eu). External generation is preferable when structured content, translation, and release automation already exist, but it creates responsibility for profile selection, identifiers, XHTML, validation, and import regression testing.

At go-live, submission remains **voluntary** for centrally authorised medicines until revised pharmaceutical legislation is fully applicable (www.ema.europa.eu). ePI does **not** replace Word or PDF product-information annexes in eCTD (www.ema.europa.eu). Readiness therefore means dual-running the established submission and the new structured lifecycle. The best near-term program freezes a versioned profile package, establishes identifier ownership, proves one-language import, runs QRD and FHIR validation locally, rehearses procedure-specific status changes, and keeps the eventual eSubmission Delivery or IRIS EPI ID declaration behind configuration because EMA says those details will be updated (www.ema.europa.eu).

Introduction and Background

Electronic product information is authorised statutory medicines information adapted for electronic handling and distribution. It covers the summary of product characteristics (SmPC), Annex II for centrally authorised products, labelling, and the package leaflet. The final guide provides validated examples for all four document classes (epi.ema.europa.eu). Sweden's regulator captures the central design objective neatly: FHIR-adapted ePI is readable by people and machines (www.lakemedelsverket.se).

That dual purpose changes the implementation problem. A team is not merely placing approved text into XML. It is maintaining regulated narrative inside a versioned resource graph whose identifiers, section codes, authorisation context, language metadata, procedure state, and presentation controls must remain aligned. The EU common standard uses Fast Healthcare Interoperability Resources (FHIR) (www.ema.europa.eu). FHIR resources can be represented in XML, JSON, and Turtle (hl7.org), with formal FHIR media types for XML and JSON (hl7.org). A FHIR document can gather its complete content into a single XML or JSON representation (hl7.org) (hl7.org), but a valid serialization is only one layer of readiness.

The audience for this guide includes regulatory operations, product-information owners, Regulatory Information Management (RIM) architects, labelling and translation leads, and integration engineers. Their immediate decision is whether to author in the PLM Portal or **generate conformant FHIR externally**. Their broader responsibility is to preserve the approved content and its lifecycle across PLM, Product Management Service (PMS), eCTD, eSubmission Delivery, IRIS, translation tooling, and internal repositories.

For an adjacent consultancy such as IntuitionLabs, the defensible role is architecture, integration, and governed workflow support, not acting as an ePI submission product. Its first-party description emphasizes data pipelines, integration, warehousing, and business intelligence (intuitionlabs.ai). That perspective is relevant to readiness design, but it should not be confused with EMA's authoritative requirements or with the PLM Portal itself.

Key Changes in the Final ePI FHIR v1.0 Guide

A released baseline, separate from operational go-live

Version **1.0.0** is a release, not another preview. EMA's releases page calls it the first release of the EMRN resource definitions and profiles (epi.ema.europa.eu). It pins the guide to package EUePI#1.0.0 and FHIR **5.0.0** (epi.ema.europa.eu). For engineering, this supplies an identifiable package to put under configuration control. Whole-version package pinning is consistent with HL7's dependency guidance (fhir.hl7.org) and permits repeatable profile validation (hl7.org).

Operationally, however, the applicant guidance states that it becomes applicable only at ePI go-live (plm-portal.ema.europa.eu). No official page reviewed for this report confirmed that production go-live had occurred by **19 September 2026**. Teams should therefore separate two backlogs:

- **Current preparation:** profile mapping, local validation, identity reconciliation, translation pilots, import tests, and operating procedures.
- **Go-live-controlled activity:** formal submission status changes, declaration through the named regulatory channel, and production publication expectations.
- **Configuration watch:** the go-live date, eligible procedures, registration access, endpoints, and EPI ID declaration fields.
- **Evidence retention:** the exact guide package, value sets, QRD material, and portal guidance used for every test run.

This distinction prevents a common planning error: treating a final data standard as proof that every surrounding production service and operating instruction is final.

The resource graph replaces document-only thinking

The final guide constrains four FHIR resource roles. The **List** is the master-level record and requires a persistent identifier (epi-dev.ema.europa.eu). A **Bundle** represents each individual document and has type=document (epi.ema.europa.eu). The first Bundle entry is a **Composition**, consistent with the FHIR document model (hl7.org) (hl7.org). The Composition carries headings, subheadings, and narrative (epi.ema.europa.eu). Images are contained **Binary** resources rather than unmanaged file references (epi-dev.ema.europa.eu). This graph keeps document content portable across the formal XML and JSON forms (hl7.org).

Table 1 maps the resource graph to regulated content and practical ownership.

FHIR layer	Regulated content or function	Primary control owner	Implementation control
List	Master grouping and lifecycle identity across ePI documents	RIM and regulatory operations	Maintain the persistent ePI master identifier and relationships to versions.
Bundle	One SmPC, Annex II, labelling, or package-leaflet document	Integration engineering	Enforce document type, complete entries, stable Bundle identity, and serialization rules.

FHIR layer	Regulated content or function	Primary control owner	Implementation control
Composition	Document metadata, QRD hierarchy, headings, narrative XHTML, language, and identifiers	Product information and labelling	Apply the general, authorisation-type, document-type, and QRD template profiles.
Binary	Images embedded in the regulated document	Content operations and accessibility	Package images as contained resources and supply meaningful alternative text, which PLM import guidance expects (plm-portal.ema.europa.eu).

The table shows why ownership cannot sit solely with an integration team. Engineering can guarantee the graph and serialization. It cannot determine approved text, the correct procedure association, or the authoritative QRD structure without regulatory and labelling owners.

CAP and NAP profiles are related, not interchangeable

The guide supplies a Composition profile for centrally authorised products (CAPs) (epi.ema.europa.eu) and another for nationally authorised products, including mutual-recognition and decentralised contexts (epi.ema.europa.eu). Their section vocabularies differ. The CAP value set contains **134 concepts** (epi-dev.ema.europa.eu); NAP-MRP-DCP contains **122** (epi-dev.ema.europa.eu). Document-specific sets further constrain the hierarchy: SmPC has **62**, labelling **34**, package leaflet **30**, and the document-type set **9 concepts** (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu).

The practical rule is to derive profile selection from explicit metadata, never from document wording. Local validation guidance says selection combines authorisation type, document type, and language (plm-portal.ema.europa.eu). It also says only one of the document-type Composition profiles should be added (plm-portal.ema.europa.eu). CAP readiness for the voluntary centralised rollout should not be delayed by an enterprise-wide NAP rollout, but a shared content model should avoid hard-coding CAP assumptions into every layer.

Implementation Considerations and Process Changes

Choose PLM authoring or external FHIR generation deliberately

EMA documents both approaches. Applicants can generate FHIR with their own systems or providers and import it (plm-portal.ema.europa.eu). They can also create and edit content in PLM. The choice should follow content scale, reuse, translation volume, validation maturity, and integration capability, not a generic preference for portals or automation.

Table 2 compares the two operating models.

Decision factor	Author in PLM Portal	Generate externally and import FHIR
Best fit	Limited initial volume, manual operating model, and a team learning the target structure	Structured content already exists, many products or languages, or repeatable automation is required

Decision factor	Author in PLM Portal	Generate externally and import FHIR
Authoring behavior	Direct use of portal controls; pasted Word styling is removed automatically (plm-portal.ema.europa.eu)	Source system must produce compliant XHTML, XML, profiles, codes, and identifiers
Validation boundary	Portal guides the author, but governance still owns approved content and metadata	Local validator and profile package become formal release controls; downloadable material includes profiles, code systems, and value sets (plm-portal.ema.europa.eu)
Translation	Portal-centered human workflow	Export, translate only permitted content, preserve structure, validate, and import one language at a time
Change control	Portal status and version workflow are the operational record	Dual change control across source, generated artifact, import result, and portal lifecycle
Strategic tradeoff	Faster initial adoption, more manual scaling	Higher engineering cost, stronger reuse and repeatability potential

The comparison is not binary for all time. A sensible phased program can author a small CAP cohort in PLM to learn the regulatory workflow, while an external generator proves conformance against the same approved examples. This extends the controlled-real-case approach used in EU-supported ePI work (health.ec.europa.eu). The decision can then be revisited using observed error rates, cycle time, and translation effort.

Make identity and lifecycle ownership explicit

Several identifiers coexist and must not be collapsed:

- **Marketing authorisation identity:** the regulated product context, mastered through regulatory data.
- **PMS identity:** the product identifier carried through the Composition extension; EMA describes PMS data as identifying a human medicinal product from regulated information (www.ema.europa.eu).
- **EPI ID:** the portal lifecycle key for the ePI version submitted in connection with procedures.
- **List persistent identifier:** the durable master-level key in the FHIR graph.
- **Bundle identifier:** a document-instance key used with language and EPI ID during import matching.
- **Composition identifier:** a key for one specific Composition version, not the whole lifecycle (epi-dev.ema.europa.eu).
- **Procedure number:** the regulatory event association needed for publication.
- **QRD version:** the structural template baseline that must match across ePI and conventional product information.

PLM import checks **EPI ID, language code, and Bundle identifier** when deciding whether an import creates or replaces a document (plm-portal.ema.europa.eu). Meanwhile, FHIR's logical ID is unique only within the same resource-type namespace on a server (hl7.org), and its version ID changes when content changes (hl7.org). An enterprise crosswalk should therefore store system, namespace, value, object type, effective dates, and lifecycle status. The broader medicinal-product identity model spans five ISO IDMP standards (www.iso.org). Never use one convenient UUID as a substitute for all business identities.

Design the procedure workflow as a state machine

At go-live, an ePI must be linked to an authorised medicinal product before it can reach Submitted status (www.ema.europa.eu). Changing status to Submitted formally sends the ePI to EMA (www.ema.europa.eu). Each procedure-linked update creates a new version with a new EPI ID (www.ema.europa.eu). Archived ePI remains stored and API-accessible but cannot be edited (plm-portal.ema.europa.eu).

Table 3 translates the September guidance into a design matrix. These are **post-go-live target rules**, not evidence that production processing is already open.

Procedure context	Create or update point	PLM status action	Declaration channel named by guidance
Initial MAA	Complete the first ePI by initial eCTD submission	Keep Draft initially (www.ema.europa.eu); after favourable outcome, update and submit with final translations	eSubmission Delivery
Type IA or IAIN; Type IB without linguistic review; Article 61(3) without linguistic review; transfer	At the same time as the procedure application (www.ema.europa.eu)	Create first ePI or update the published ePI, then move to Submitted	eSubmission Delivery
Type IB with linguistic review; Article 61(3) with linguistic review; other post-authorisation procedures	With final product-information translations (www.ema.europa.eu)	Update, validate, associate the procedure, then submit	IRIS Industry Portal
Published ePI superseded by a later version	After the procedure-linked replacement is published	Prior version becomes Archived and non-editable	Portal and repository lifecycle

The matrix should be encoded as configurable workflow rules, not embedded in content-generation code. The guidance still says detailed instructions for communicating the EPI ID through eSubmission Delivery or IRIS will be added (www.ema.europa.eu). That unresolved field is a tracked dependency, not a reason to delay all preparation.

Data and Workflow Architecture

Coexistence with eCTD, Word, and PDF

ePI adds a structured lifecycle. It does not remove the established electronic Common Technical Document (eCTD) lifecycle. EMA explicitly says ePI does not negate the requirement to submit product-information annexes (www.ema.europa.eu). The ePI and the Word or PDF product information must use the same valid QRD template version (www.ema.europa.eu). EU Module 1 carries language, document-type, and country identifiers for product information (esubmission.ema.europa.eu), and its XML backbone is validated against the EU regional DTD (esubmission.ema.europa.eu). ICH maintains a separate eCTD release track (admin.ich.org).

This creates a controlled dual representation. The approved narrative must remain semantically identical while its structures differ. Reconciliation should compare at least:

- **Document inventory:** all expected SmPC, Annex II, labelling, and package-leaflet documents exist.
- **Template baseline:** the same QRD version is recorded and applied.
- **Language set:** the same intended languages and market context are represented.
- **Section sequence:** mandatory and optional sections remain in valid positions.
- **Text content:** approved narrative, tables, lists, symbols, and references reconcile.
- **Procedure association:** the eCTD sequence and ePI version point to the same regulatory event.
- **Approval state:** draft, submitted, approved, published, and archived states do not drift.

Working documents in an eCTD package belong in a separately named xxxx-workingdocuments folder whose number matches the sequence ([esubmission.ema.europa.eu](https://www.esubmission.ema.europa.eu)). The backbone remains subject to the EU regional DTD ([esubmission.ema.europa.eu](https://www.esubmission.ema.europa.eu)), while eCTD standards continue on their own ICH schedule (admin.ich.org). That packaging rule should remain in the submission runbook alongside, not inside, the ePI generation pipeline.

Translation, styling, and accessibility

During the voluntary phase, EU-language translations, including Icelandic and Norwegian, are optional (www.ema.europa.eu). EMA expects all-language submission to become mandatory later (www.ema.europa.eu). EU law already requires labelling particulars in the official language or languages of the relevant Member State (eur-lex.europa.eu), and EU-supported ePI work explicitly includes linguistic review and translation upload (health.ec.europa.eu). Systems should be designed for the future language cardinality even if the first voluntary cohort is smaller.

The safe translation pattern is controlled extraction and reinsertion. EMA instructs translators to change only indicated parts so the FHIR XML structure is preserved (plm-portal.ema.europa.eu). PLM currently imports only one language at a time (plm-portal.ema.europa.eu). Language remains a statutory market attribute (eur-lex.europa.eu) and an accessibility input for multilingual speech synthesis (www.w3.org). Consequently, a translation manifest should bind source hash, target language, translator package, Bundle identifier, EPI ID, QRD version, validation report, and import outcome.

Styling also requires separation of concerns:

- **Narrative format:** section content is FHIR narrative XHTML (plm-portal.ema.europa.eu).
- **Escaping:** special symbols must be escaped so XHTML, XML, or JSON is not broken (plm-portal.ema.europa.eu).
- **Presentation:** exported FHIR XML does not contain QRD styles (plm-portal.ema.europa.eu).
- **Language metadata:** FHIR language tags support indexing and text-to-speech accessibility (hl7.org).
- **Web accessibility:** WCAG 2.2 requires a page's default human language to be programmatically determinable (www.w3.org) and supports language-aware speech synthesis (www.w3.org).

Reference architecture and phased test plan

A practical reference architecture has six separable services: source content and approval, identifier registry, FHIR transformation, local validation, PLM interaction, and evidence archive. At an API boundary, a FHIR CapabilityStatement can describe whether a system initiates or receives RESTful operations (fhir.hl7.org) and expose client-relevant security information (fhir.hl7.org). Treat the currently documented public API carefully. EMA's API page describes keyless access to pilot-published ePI (epi.developer.ema.europa.eu); it does not establish a production upload API, authentication scheme, service level, or \validate support.

The phased test plan should be:

1. **Package test:** pin EUePI#1.0.0, its dependencies, value sets, and QRD artifacts. HL7 recommends whole package version numbers for dependencies (fhir.hl7.org).
2. **Golden-example test:** parse EMA examples for each document type and compare the generated graph across supported serializations (hl7.org).
3. **Profile test:** validate Bundle, general Composition, CAP or NAP, document type, QRD template, terminology, and narrative (hl7.org).
4. **Negative test:** mutate identifiers, missing sections, invalid codes, broken XHTML, language, and Binary references; assert structured diagnostics (fhir.hl7.org).
5. **Round-trip test:** export, import, re-export, and compare normalized resources without expecting presentation styling in XML. Resource version IDs should change when content changes (hl7.org).
6. **Translation test:** extract permitted strings, translate one language, reinsert, validate, import, and reconcile rendering.
7. **Lifecycle test:** rehearse first creation, update, Submitted, published, superseded, and Archived states with synthetic procedure numbers.
8. **Interface test:** confirm actual endpoint capabilities before use. FHIR does not require all servers to support \validate (fhir.hl7.org).
9. **Security test:** apply the published interface security model; HL7 recommends TLS for FHIR transport (hl7.org), and if OAuth bearer tokens are used, TLS is mandatory under RFC 6750 (www.rfc-editor.org).
10. **Release evidence test:** archive source hash, generated artifacts, validator version, profile package, errors, approvals, portal receipt, and reconciliation.

Local validation should return structured diagnostics. FHIR defines OperationOutcome as the carrier for issues found by a validation service (fhir.hl7.org) and profile checking as a distinct validation concern (hl7.org). Even a successful \validate response is not a guarantee that a later create or update will succeed (hl7.org). Portal import therefore remains a distinct acceptance test, not a substitute for local conformance (fhir.hl7.org).

Data Analysis and Evidence

The evidence supports readiness investment, but it does not support assuming a fully scaled production service. EMA's pilot ran for **one year**, from July 2023 through August 2024 (www.ema.europa.eu). It included **23 regulatory procedures** (www.ema.europa.eu). Participants created, submitted, and updated ePI for **92%**, or **23 of 25**, co-opted procedures (www.ema.europa.eu). EMA also reported publication for **100%** of submitted ePIs with a positive outcome (www.ema.europa.eu).

The average industry creation time in the pilot was **5.2 hours** (www.ema.europa.eu). That is useful as a directional baseline, not a universal business case. The sample reflects pilot participants and procedures, and it does not isolate source quality, document complexity, language count, rework, or enterprise integration cost. Separately, **20 organizations** completed FHIR Import user acceptance testing (plm-portal.ema.europa.eu). That demonstrates tested import capability, not production availability for every applicant.

The guide's vocabulary sizes indicate why validation must be data-driven. CAP has **134** section concepts versus **122** for NAP-MRP-DCP; SmPC alone has **62**, labelling **34**, and package leaflet **30** (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu) (epi-dev.ema.europa.eu). Hand-maintained mappings will be fragile unless codes, effective versions, and source artifacts are imported into governed reference data.

A reader-supplied readiness score

No public evidence supports assigning a universal readiness percentage to an organization. Teams can calculate their own score using only observed control completion: **dimension score = completed applicable controls / applicable controls x 100**. Exclude non-applicable controls from both numerator and denominator. Report four dimensions separately, then calculate an unweighted overall score only if all four have at least one applicable control. The controls deliberately reflect FHIR profile validation (hl7.org), IDMP's multi-standard identity scope (www.iso.org), and programmatically determinable language (www.w3.org).

Data controls

- **Identity crosswalk complete:** authorisation, PMS, EPI, List, Bundle, Composition, language, and procedure keys have owners.
- **Profile rules configured:** CAP or NAP, document type, language, and QRD version drive selection.
- **Terminology versioned:** code systems and value sets are imported and traceable to package 1.0.0.
- **Narrative controlled:** XHTML, tables, lists, symbols, and Binary references have transformation tests.
- **Approved-text reconciliation defined:** structured output can be compared with Word and PDF.
- **Retention defined:** source, output, validation, approval, and receipt evidence are immutable and searchable.

Workflow controls

- **Procedure matrix approved:** creation, update, timing, status, and channel are encoded for applicable procedures.
- **Dual submission rehearsed:** ePI and existing eCTD annexes remain coordinated.
- **Status authority assigned:** only authorized roles can move Draft to Submitted.
- **Version transition tested:** published records can be superseded and prior versions archived.
- **Go-live dependency tracked:** current preparation is separated from future production steps.
- **EPI ID declaration configurable:** no fixed field is assumed before EMA updates its instruction.

Translation controls

- **Language manifest implemented:** every artifact binds language, identifiers, hash, and QRD version.
- **Protected structure enforced:** translators modify only approved narrative segments.
- **One-language import tested:** the current PLM constraint is reflected in orchestration.

- **Accessibility checked:** language tags, image alternative text, reading order, and output rendering pass review.
- **Linguistic review aligned:** ePI timing matches final product-information translation timing.
- **Scale model tested:** future all-language submission does not require redesign.

Integration controls

- **Package pinned:** validator and EUePI#1.0.0 dependencies are reproducible.
- **Golden files automated:** all four document types have positive regression tests.
- **Negative tests automated:** invalid codes, identifiers, sections, and XHTML produce expected findings.
- **Portal round trip completed:** import, export, normalization, and reconciliation succeed.
- **Capability discovery documented:** production endpoints, operations, security, and limits are verified rather than assumed.
- **Monitoring implemented:** failures are correlated by product, EPI ID, Bundle, language, procedure, and release.

This score is intentionally evidence-based. For example, completing four of six applicable integration controls yields **66.7%**, but the result should not be averaged away if the missing controls are portal round trip and capability discovery. A dashboard should show both percentage and open control names.

Implications and Future Directions

The final guide moves ePI from exploratory mapping toward implementable configuration. The next constraint is operational certainty. The draft roadmap anticipates phased CAP go-live from **Q4 2026** (www.ema.europa.eu), while the submission guidance remains explicitly prospective. Organizations should monitor official change points, not infer them from a technically active guide website.

Future design should also anticipate broader interoperability. EMA is implementing Identification of Medicinal Products (IDMP) through a phased Substance, Product, Organisation and Referentials (SPOR) program (www.ema.europa.eu). The ISO IDMP family contains five core standards (www.iso.org). European Commission funding material calls for interoperability with Union and global initiatives such as the European Health Data Space (health.ec.europa.eu) and previously funded controlled real-case evaluation (health.ec.europa.eu). The human-readable and machine-readable objective remains central (www.lakemedelsverket.se). These signals favor explicit identifiers, standards-based interfaces, and loose coupling over portal-specific data duplication.

The regulatory and technical release calendars must remain independent. ICH endorsed **eCTD v4.0 Implementation Guide 1.7** and Controlled Vocabulary Package **1.0.3** in June 2026 (admin.ich.org). EU Module 1 separately identifies product-information language, document type, and country (esubmission.ema.europa.eu). That change does not eliminate the parallel ePI lifecycle. A durable architecture maps both to authoritative content and regulatory events without forcing one standard's identifiers or release schedule into the other.

IntuitionLabs' adjacent perspective belongs in this orchestration layer. Its own service description includes strategic guidance on digital transformation and technology roadmapping (intuitionlabs.ai). For a regulated ePI program, such advisory work should be measured by traceability, reproducible validation, and successful controlled handoffs, while EMA and the standards bodies remain the sources of requirements.

Frequently Asked Questions (FAQs)

Is EMA ePI FHIR 1.0.0 final?

Yes, as a technical implementation-guide release. The releases page identifies **1.0.0** as the first released EMRN resource-definition and profile package on **17 September 2026** ([epi.ema.europa.eu](https://www.ema.europa.eu/ema/press/news/2026/09/17/ema-releases-1-0-0-version-of-ema-e-pi-fhir)). Final technical publication is separate from production go-live.

How does an applicant submit ePI to EMA today?

As of **19 September 2026**, organizations should prepare and test but verify that go-live has occurred before following the prospective production workflow. EMA states that the centralised-procedure guidance becomes applicable at go-live (plm-portal.ema.europa.eu). Once applicable, the ePI is created or imported in PLM, linked to the authorised product, validated, associated with procedures, moved to Submitted at the applicable time, and its EPI ID declared through eSubmission Delivery or IRIS. The exact declaration mechanics are still pending.

Does ePI replace Word or PDF product information in eCTD?

No. EMA explicitly retains the product-information annex requirement (www.ema.europa.eu). The structured and conventional versions must remain synchronized, including their QRD template version.

Can a company generate ePI outside the PLM Portal?

Yes. EMA says applicants can generate FHIR ePI using their own systems or providers and import it (plm-portal.ema.europa.eu). External generation makes local conformance, terminology, identity, XHTML, translation, and regression controls the applicant's responsibility.

Are translations mandatory during the voluntary phase?

No. EU-language translations, including Icelandic and Norwegian, are optional in the voluntary phase (www.ema.europa.eu). However, EMA expects all-language submission to become mandatory, so architecture should scale beyond the initial cohort.

Is there an EMA ePI production API for submission?

The reviewed official material does not document a production upload API, supported FHIR interactions, authentication scheme, or service levels. The public consuming API page describes keyless access to pilot-published content (epi.developer.ema.europa.eu). The verified applicant path is PLM authoring or FHIR XML import. System designs should keep API integration modular until EMA publishes production capabilities.

Conclusion

EMA ePI FHIR v1.0.0 is ready to serve as a controlled engineering baseline. It defines a coherent resource graph (epi.ema.europa.eu), separate CAP and NAP constraints (epi.ema.europa.eu), document and section vocabularies, identifier extensions, narrative rules, examples, and downloadable validation material. Those controls align with

profile-based FHIR validation (hl7.org), multi-standard medicinal-product identity (www.iso.org), and language accessibility (www.w3.org). The release is not evidence that every production workflow, endpoint, registration route, or declaration field is already available.

The implementation priority is therefore disciplined preparation. Regulatory operations should own procedure and status rules. Product-information teams should own approved content and QRD structure. RIM should own cross-system identity and lifecycle reconciliation. Translation leads should protect structure and language metadata. Engineering should pin the package, automate positive and negative validation, prove one-language imports, and retain release evidence.

For the voluntary centralised-procedure rollout, the safest decision framework is pragmatic. Use PLM authoring when volume is limited and learning speed matters. Use external FHIR generation when structured reuse, multilingual scale, and automation justify the additional validation and integration controls. In either model, retain the existing eCTD Word and PDF annex process, because ePI complements rather than replaces it.

Finally, maintain a clear boundary between what is final and what is prospective. The **1.0.0 data standard is final**. The September workflow guidance is tied to go-live. EPI ID communication details are still pending. Treat those facts as separate configuration states, measure readiness with completed controls rather than optimism, and the organization can move from pilot learning to production without rebuilding its information architecture at each regulatory update.

For informational purposes. Verify critical medical, legal, financial, regulatory, and technical information with qualified professionals and primary sources.