

FUTURE-PROOFING GOVERNED
CONTENT ACROSS THE ENTERPRISE

One Drug, One Truth

What it costs when your approved content is maintained in scattered copies instead of one governed set of core data and metadata

Executive Summary

Somewhere in your organization, each individual drug's approved dose statement is being maintained as scattered copies, typically kept aligned by hand, across multiple formats. That is the shape of a liability most pharmaceutical companies carry without naming it. When a drug has no single governed set of core data and metadata, the same approved truth, one dose, one safety statement, one manufacturing commitment, is maintained separately everywhere it appears. It lives as a human-readable document. It now must also live as several machine-readable renderings, the different formats each regulator requires. Keeping every copy identical, and provably so, is no longer tenable by hand.

This obligation was never designed. It accreted one mandate at a time. Electronic product information (ePI in the EU, SPL in the US), the harmonized clinical protocol (ICH M11), structured manufacturing data (PQ-CMC, with ICH M4Q(R2) as its international counterpart), and the newest submission rail (eCTD v4.0) each added a requirement and retired nothing. So the work of keeping one truth aligned grows with every product, market, language, and format. It grows fastest after approval, where most regulatory effort lives. Quietly, it also caps the return on the structured-content and AI investments a company has already made.

Few leaders can say with confidence what this costs, because no one owns it.

It sits inside regulatory operations, stretches across systems bought separately, and appears on no report a senior leader routinely sees. More bolt-on tools will not solve it.

The way out is structural: a single governed source of machine-readable content from which every required format is generated.

One approved truth is authored once and rendered everywhere, with its full footprint visible before anything changes. Artificial intelligence does not change that logic. AI amplifies whatever content it is given, which makes the governed source the precondition for using it safely, and the lever that turns it from a drafting aid into a way to cut the costliest work.

The same governed source that keeps every rendering aligned is also what makes the next thing possible: product information that can be focused to an individual patient's needs, not just filed with a regulator.

This liability scales with portfolio size, but it is not confined to the largest companies. Any sponsor, generics and mid-size manufacturers included, that maintains approved content across multiple markets and formats carries a version of it.

This paper does three things. It shows how to see the liability and size it against your own portfolio. It explains why a decade of investment has not compounded. And it describes where to focus funding and effort. The argument rests on the industry's own regulatory benchmark and on primary regulatory sources.

A real-world example:

The statement that lives in six places at once

Take a single approved sentence, a dose statement for one medicine, and follow it across the systems that must carry it. The same sentence must appear, ***unchanged in meaning***, in the US prescribing information, in the EU product information, in the summary of product characteristics and patient leaflet, inside the submission dossier, and in any medical response that quotes the label. Each destination encodes it differently.

One approved sentence, at least six renderings:

Rendering	Region	Standard
Prescribing information (USPI)	United States	SPL XML/FHIR
Product Monograph	Canada	XML-PM
Electronic product information (ePI)	European Union	FHIR
Summary of product characteristics and patient leaflet	European Union	ePI and Patient Leaflet
Submission dossier	All ICH regions	eCTD
Medical information response	All markets	Queries to HCPs

A change to the source must reach all of them and read identically in each.

This did not happen by accident, and it did not happen at once. Over two decades, regulators layered structured, machine-readable representations on top of the documents they already required. They removed none of the documents. The Word and PDF dossier still governs assessment, and the structured renderings sit alongside it. The obligation is additive by design.

Two levels are worth separating, because they call for different things, and the distinction governs what you buy. The first is the transport rail, the way a whole dossier is packaged and exchanged with an agency. That rail is the electronic common technical document (eCTD). Even its newest version, v4.0, runs on the Health Level Seven (HL7) version 3 Regulated Product Submission (RPS) standard rather than on Fast Healthcare Interoperability Resources (FHIR).⁸ The second level is the content itself, rendered for each destination, and here there is no single format. Electronic product information (ePI) and structured manufacturing data use FHIR. The harmonized clinical protocol uses a FHIR-based exchange. US labeling as Structured Product Labeling (SPL) is XML today and moving to FHIR, while safety case reports use HL7 version 3.¹¹ Clinical study design is captured in the CDISC Unified Study Definitions Model (USDM), exchanged via an API in JSON or XML and interoperable with FHIR.¹² Product identity is carried in the ISO standards known as Identification of Medicinal Products (IDMP).⁷ The renderings are heterogeneous. A platform built around any single format cannot produce the full set on its own, and that single fact calls into question many of the tools companies have already bought.



Numbers Game

The exposure becomes concrete the moment you size it, and the inputs are numbers a company already holds:

Active products

- × markets per product
- × languages per market
- × renderings per market
- × lifecycle changes per year

**the number of times
one truth must be kept
aligned**

For any global portfolio that number is large, and most organizations have never counted it. Language alone multiplies it fast: a single EU centralized approval requires the product information in up to 24 official languages, and EU law requires every language version to be identical in content, apart from national particulars such as the member-state 'blue box'. That uncounted number is the liability. It scales with the business, and today it lives inside regulatory operations, with no enterprise owner and no line on any executive dashboard.

The same liability, three shapes

The formula holds for every company that carries approved content, but its weight lands differently depending on the shape of the business.

A large originator feels it as breadth and depth at once: many products, each in many markets and languages, each anchored on a company core data sheet that must stay synchronized with every downstream label.

A generics or biosimilar company feels this differently but no less. It rarely holds a company core data sheet, because it did not originate the molecule — its labels are driven by the originator's. When that reference label changes, for any reason, the change has to be absorbed and pushed across the company's own regional labels and localizations, on the agency's clock. The difficulty is volume. A portfolio measured in hundreds or thousands of SKUs and market combinations turns each externally-

triggered change into a wave of manual reconciliation. For a company of this shape, the governed source is a throughput proposition: propagate an inbound label change across the whole portfolio quickly and provably, instead of one SKU at a time.

A small or emerging biotech sits at the other end, with a handful of marketed products or none yet. It is tempting to treat governed, structured content as a problem for later, at scale. The opposite is true. The concepts return the same value at any size and adopting them early means the discipline is built in before the copies multiply, with no decade of accreted document-and-copy practice to unwind. It also compounds through growth: as the company expands and acquires, a content strategy set from day one makes each integration a transfer rather than a reconciliation. Starting governed is the cheapest this will ever be.

This is not master data management. Master data, including IDMP and SPOR, governs how a product and its substances are identified and classified. One drug, one truth governs something different: the approved statement itself, the dose, the contraindication, the established condition, with its applicability and lineage tracked across every market and format. Master data answers which product this is.

A governed statement answers what was approved to be said about it, and proves it stayed the same everywhere.

Why the gap widens every year

The exposure would be manageable if it held still. It does not. It grows, and the growth accelerates as the business does, for three reasons that compound one another.

More mandates, each one additive

The mandates keep arriving, and each new one adds a rendering without retiring an old one. ePI is phasing in from 2026 and becomes mandatory for newly authorized medicines under the revised EU pharmaceutical legislation.¹ The harmonized clinical protocol standard, ICH M11, reached its final step in November 2025.⁵ Structured manufacturing data is arriving in phases. The v4.0 submission rail is voluntary in the United States today, mandatory in Japan as of 2026, and expected in the EU around 2027. Every one is another form the same approved truth must take on top of the forms already required.

Review is turning machine-readable

Review itself is turning machine-readable, which changes what inconsistency costs. The FDA is building the ability to compare structured content across a submission. Two programs drive this: its Knowledge-Aided Assessment and Structured Application (KASA) framework, and its PQ-CMC program for structured manufacturing data.⁶ When a reviewer can compare structured elements, an internal contradiction that a human might once have read past instead surfaces as a data-level discrepancy. Tolerance for that kind of divergence falls. The direction shows in the regulators' own tools, limits included. The FDA has deployed an internal assistant that can compare labels and summarize safety reports. The same tool has produced fabricated citations, which is exactly why it has been kept out of formal assessment.² The signal is not that machine-assisted scrutiny is finished. It is that it is coming, and that it rewards companies whose renderings already agree.

The burden lands hardest after approval

Most regulatory effort across a portfolio is spent

after approval, not at first submission. There, a single safety change never moves through just one document. The change must reach labeling in every market, update both the FHIR and SPL renderings of that label, align with medical information, and flow into the safety and regulatory records. Each path runs on its own clock and in its own format.

In the EU, that clock is formal. A safety change to a centrally authorized product re-enters a standardized linguistic-review process that runs about 25 days across the EU's 24 official languages plus EEA languages, and each member state can return national comments before the label is final. Other regions run their own cycles on their own timing.

Today that reconciliation is held together by people working by hand, and manual coordination at that scale has a cost.

It shows up in forms a senior leader recognizes at once:

- Regulator queries that delay approvals.
- Late label changes that push back a launch or a market action.
- Effort duplicated across functions and markets.
- Inspection and audit exposure when two renderings of the same truth quietly diverge.
- Each fresh mandate arriving as an unbudgeted fire drill rather than a routine output.

Suppose a routine safety variation updates one warning. In a single-source operation, the warning is changed once and every dependent rendering and language version regenerates from it. In a document-and-copy operation, the same change spawns dozens of separate edits, each translated, each reconciled

against the approved English, each tracked by hand, and any one of them can drift. The work is proving, across every market and language, that the same decision now reads the same everywhere.

None of this is overstated. The shift is real but still phased and partly voluntary, the human-readable document still governs assessment, and the largest value figures in circulation are estimates rather than measured results. The trajectory, though, is not in question.

The same gap blocks the patient

There is a further cost that never reaches a regulatory dashboard: the patient never sees the benefit of content the company cannot keep coherent.

The industry's direction of travel is clear. The Gravitare Health initiative, a public-private partnership co-led by the University of Oslo and Pfizer, built a mechanism (the G-lens®) that focuses approved electronic product information to a patient's own context, language, and health literacy without hiding or altering the approved text, running on FHIR-based ePI. That is the point: a focusing layer can only work on content that is already governed, structured, and machine-readable at the source. A company whose ePI is maintained as scattered copies cannot feed it reliably, however good the patient-facing tool becomes. The fragmented source that raises compliance cost is the same thing that keeps personalized, patient-readable information out of reach.¹⁷



Why the spending hasn't paid back



What the benchmark shows

Gens & Associates' 2025 World Class RIM study examined 59 life sciences organizations, and its findings are sobering for anyone who assumes their investment has bought readiness. Only one of the 59 was judged genuinely prepared for the disruption ahead. What

sets that organization apart, in the study's own analysis, is not its technology or its choice of vendor but the layer beneath: cross-functional collaboration, mature processes, and above all how its data is governed and owned.⁴

The study is direct on the point that readiness has little to do with AI. The differentiator is governance and accountability over the data itself. That reframing matters here, because accountability in regulatory work has traditionally sat at the level of the dossier and moving it down to the level of the content and data is what separates the prepared from the rest.

Most companies have already spent against this problem. For most, the spending has not returned what leadership expected. **The reason is not a lack of effort, and it is not the wrong vendor. It is architectural, and the industry's own benchmark says so.**

Spending conditions overlap

Most organizations are earlier than their spending suggests, and they tend to sit in several conditions at once rather than one tidy category. The conditions overlap, and a single company will often recognize itself in more than one:



Authoring in documents, reconciled by hand.

This remains the operative norm in manufacturing content, where no current mandate compels structuring, and dossiers are assembled manually under existing procedures.¹⁴



Systems modernized, content not governed.

Roughly 80 percent of top pharma are modernizing regulatory information systems, by one 2025 benchmark.³ This is why a modern RIM or publishing platform does not, on its own, close the gap: it manages the dossier and its versions, but still treats the dossier as the unit of truth, so the renderings underneath it keep their own copies.



Local structured-content wins that do not compound.

One top-five company reported a 15 percent productivity gain in approval-linked work after adopting structured authoring at the core data sheet. Such wins tend to stall at the boundary of the function that achieved them, because the source was never shared across functions.



AI pointed at existing content.

Proof-of-concept results look impressive, but practitioners warn that the appeal pulls organizations away from the harder prerequisite, structuring the content first, since a model can only populate a foundation that has been built.¹⁶

Why none of it compounded

Underneath all four conditions sits one cause, and it is structural. Money spent inside that arrangement cannot compound, because the arrangement itself is what multiplies the work. That is why readiness tracks governance over content, rather than the size of the technology budget.

AI sharpens the point rather than softening it. Generative assistance does cut authoring time. One program reduced first-draft clinical study report time from roughly 180 hours to 80, and halved errors.³

On a coherent source, that speed compounds into advantage. On an incoherent source, it manufactures divergence faster. The point is not that AI is premature. It is that AI amplifies the content beneath it, for better or worse, which is why the order of operations, governed source first, decides whether speed becomes advantage or divergence.

Whatever has been funded so far, the return is capped until governance reaches the content itself. Another increment of the same spending will not lift that cap, because the constraint is structural.

What a governed source changes

If more of the same spending will not work, **the question is what will.**

There is a way to carry this obligation without the manual reconciliation that makes it so costly, and it is structural rather than incremental. Govern the content at its source, and generate every rendering from that one governed source, rather than maintaining each rendering on its own. Three properties define what that requires.

The first is a single governed source. Every machine-readable rendering and every document is produced from it. Nothing is authored or maintained in parallel.

The second is that content and its governing metadata travel together. The content is the approved truth itself: product and substance identity, indication and dosage, contraindications, warnings, and adverse reactions, efficacy claims and composition, and the manufacturing commitments that ICH Q12 designates as established conditions, where a change carries a regulatory obligation.¹³ The metadata keeps that truth traceable as it moves. It carries the product and substance identifiers; the market, indication, and claim type; the approval status and version; the language link; the lineage of every component; and the audit trail required by 21 CFR Part 11 and EU Annex 11. When the two travel together, a change to a source statement reveals every place that statement lands, across markets, document types, and formats, before anyone acts on it.

The third property concerns AI, and it deserves a clear statement, because AI is now one of the most discussed levers in pharmaceutical content. The position here is specific. AI stays bounded to the governed source, and people remain accountable for what it produces. It generates from approved content, and it does not author the truth on its own. This is the same discipline regulators now apply to themselves. The FDA's 2025 draft guidance asks sponsors to define an AI model's context of use, assess its risk, and document the result before that output supports a regulatory decision.² The principle for content is identical. AI accelerates whatever control model already sits beneath it. On governed content it becomes a force multiplier for consistency, impact analysis, and controlled reuse. On ungoverned content it accelerates output while leaving provability unchanged, and risk scales faster than the capacity to review it.

What the right AI actually does

Governed content also changes what the most useful AI actually does. The familiar pattern, retrieval-augmented generation, writes new text and then asks a person to check it. That accelerates drafting, which is rarely the expensive step. The expensive steps are reconciliation, impact assessment, and re-review, and they grow faster than headcount as reuse and variation increase. A governed source allows a different order of operations, and EU localization is the clearest case of it in practice:

1

Reuse (RARE, retrieval-augmented reuse). Reuse an approved, human-reviewed component verbatim when one fits, because it already carries its metadata, lineage, and approvals.

2

Transform (RAT, retrieval-augmented transformation). Apply a controlled, track-change-visible adaptation when a region needs one, such as a national blue-box statement or a member-state language preference applied during linguistic review, with the source provenance preserved and the reviewer able to accept or reject each edit.

3

Generate (RAG, retrieval-augmented generation). Generate genuinely new text only when no governed content exists, and even then within metadata constraints and full attribution.

Reuse first, transform second, generate last. The safest output is approved content rather than new prose, so the work that actually consumes time, human re-review, shrinks instead of growing. That inverts the usual economics of AI in regulated content, and it is available only when the content beneath the model is governed and structured to begin with.

Only one operating model can carry this liability: a governed source from which every rendering is generated. Docuvera was built to run it. It places structured authoring at the core, with governance and metadata enforced at the moment content is created, and it spans the full estate, across chemistry, manufacturing, and controls (CMC), clinical, labeling, safety, medical affairs, and quality, rather than a single function. From that one governed source, it produces every required rendering, the FHIR, HL7 version 3, JSON, and ISO outputs alike.

Its AI works inside that governance through a reuse-first hierarchy, the Hierarchy of Intelligence, which reuses approved components first, transforms them under review when a market needs an adaptation, and generates new text only as a last resort. One global regulatory team applying this model to a Clinical Overview update reused 64 percent of validated components and cut review cycle time by 27 percent, gains its quality reviewers attributed to governed reuse rather than uncontrolled automation.¹⁵



What it changes for the business

The payoff reads in the language of economics, and it is the leader's to govern:

- Timelines become predictable, because propagation is systematic.
- The marginal cost of adding a market or a format falls, because a new rendering is generated rather than rebuilt.
- A change moves quickly and safely, with its full footprint visible before it is made.
- A new mandate is absorbed from the existing source, rather than funded as another re-platforming.
- Portfolio and acquisition integration get easier, because governed, structured content transfers without manual reconciliation.
- AI finally compounds, because it runs on a foundation coherent enough to trust.
- The patient-facing future becomes reachable, because focused, personalized product information can only be generated from a governed, machine-readable source.

Where to start

The liability compounds whether or not a company acts. The most important move is the leader's to make, costs very little, and asks for none of the larger commitment up front. Commission a dependency audit of a single content type, labeling or manufacturing. Require the organization to show, for one approved statement, every place it appears across documents, markets, and systems. The result is an enterprise risk picture. It is also the real number behind the sizing logic above, the figure that turns an abstract exposure into something the leader owns.



When the map cannot be produced at all. Treat content coherence as an owned enterprise risk. Assign clear accountability for it, and fund a single cross-functional truth first, the labeling chain anchored on the company core data sheet (CCDS, or, for a company without one, the inbound reference label and its own core label set), as proof that one change can propagate and be traced across its renderings.⁹



When local wins exist but have not compounded. Stop funding isolated tools. Redirect the investment toward extending one governed truth across a function boundary, and tie further funding to cross-functional reuse rather than another local victory.

If the organization cannot produce that map, the inability is itself the finding. It means the exposure has no owner today, and naming one is a decision only a senior leader can make.

What to fund after the audit depends on what it reveals.

The situations below are conditions, not rigid stages, and an organization often shows more than one.



When structure was adopted and then stalled. Recognize the stall as structural. Move authority to the reusable statement and its established conditions, and change the measure of success from a platform deployed to reuse trusted and review narrowed to only what has changed.



When several domains are already governed. Require the single source to produce every rendering, and prove propagation across all of them. Only then scale AI on that foundation, where the investment can finally return.

One principle orders these decisions.

Fund only the moves where each step lowers the cost of absorbing the next mandate. Refuse any move whose main effect is to add one more parallel system to maintain.

Questions to ask your organization



Take these into the room. For several of them, **the inability to answer is the answer.**

1. For one approved statement, a key dose or safety warning, how many distinct renderings and locations does it occupy across our products and markets, and can our systems produce that list on demand?
2. When a single safety change is approved, how long does it take to reach every dependent rendering and market, and how do we prove that it reached all of them?
3. What share of our regulatory effort and cost is post-approval maintenance, rather than initial submission?
4. What did our prior structured-content and AI investments actually return, and why have local wins not compounded across functions?
5. Can we produce the next required machine-readable rendering, a FHIR ePI, an M11 protocol, structured manufacturing content, from our current source, or would each one be a separate build?
6. If a regulator's data-driven review compares our structured submissions, are we confident the renderings are internally consistent?

Conclusion

The obligation at the center of this paper is not optional, and it does not pause. One approved truth has to exist in every format regulators require, the number of those formats keeps growing, and the work of keeping them aligned compounds whether or not a company addresses it.

That last point is worth sitting with, because it turns inaction into a decision in its own right.

To not decide is to let the cost keep accruing, and to keep capping the return on investments already made. Liability is carried either way. The only real question is whether a company carries it deliberately, with one governed source behind every rendering, or absorbs it by default and pays for it in queries, delays, and rework it never traces back to the source.

A single source of approved truth changes the economics of carrying it. It turns a growing,

unmanaged exposure into a cost a company can predict and control. And the thing that separates the prepared from the rest, in the industry's own benchmark, is governance and accountability over the content, more than the size of any technology budget. That is a decision a senior leader is positioned to own and to make.

Docuvera exists to make that decision executable, by governing content at its source and generating every required rendering from it.

The first step toward it asks for none of that commitment. It asks only for the will to see the liability clearly. The same decision that controls cost is what eventually lets a patient read the right information, in their own language, at their own level; the governed source is the precondition for both. Commission the audit. Read the map. Decide who owns the truth.

Glossary

SCA, structured content authoring.

CMC, chemistry, manufacturing, and controls.

CCDS, company core data sheet, the master labeling document that local labels derive from.

CCSI, company core safety information, its safety subset.

SmPC, summary of product characteristics, the EU label.

USPI, United States prescribing information.

SPL, structured product labeling, the HL7 version 3 format for US labeling.

HL7, Health Level Seven, a family of health-data standards.

ePI, electronic product information.

FHIR, Fast Healthcare Interoperability Resources.

eCTD, electronic common technical document.

RPS, regulated product submission, the HL7 version 3 basis of eCTD v4.0.

IDMP, identification of medicinal products, a set of ISO standards.

SPOR, the EU substance, product, organization, and referential data services.

ICH M11 / CeSHarP, the harmonized structured clinical protocol standard.

PQ-CMC, structured pharmaceutical quality and manufacturing data.

KASA, the FDA knowledge-aided assessment and structured application framework.

ICH Q12, the lifecycle management guideline, including established conditions.

Established conditions, the elements of a dossier that constitute regulatory commitments.

Context of use, the defined scope of an AI model's role.

RAG, retrieval-augmented generation, AI that generates new text from retrieved source material.

RARe, retrieval-augmented reuse, AI that reuses an approved, human-reviewed component verbatim.

RAT, retrieval-augmented transformation, AI that adapts an approved component with visible, reviewable track changes.

ICSR / E2B(R3), the individual case safety report and its HL7 version 3 exchange standard.

ICH M4Q(R2), the ICH guideline harmonizing the quality (CMC) module of the dossier; the international counterpart to the US PQ-CMC program.

G-lens® / Gravitate Health, an IMI public-private initiative that demonstrated “focusing” of approved FHIR electronic product information to an individual patient's context, language, and health literacy.

Endnotes

1. European Medicines Agency. "Electronic product information (ePI)." ema.europa.eu/en/human-regulatory-overview/marketing-authorisation/product-information-requirements/electronic-product-information-epi. ePI FHIR Implementation Guide: [epi.ema.europa.eu/fhirig/](https://ema.europa.eu/fhirig/).
2. U.S. Food and Drug Administration. "Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products." January 2025. fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological. FDA Elsa launch, June 2, 2025: fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people.
3. McKinsey & Company. "Rewiring pharma's regulatory submissions with AI and zero-based design." 2025. mckinsey.com/industries/life-sciences/our-insights/rewiring-pharmas-regulatory-submissions-with-ai-and-zero-based-design. Value range: "Generative AI in the pharmaceutical industry." mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality. CSR result: "With gen AI, Merck and McKinsey transform clinical authoring." June 27, 2025. mckinsey.com/about-us/new-at-mckinsey-blog/with-gen-ai-merck-and-mckinsey-transform-clinical-authoring.
4. Gens & Associates. "2025 Operational Excellence and World Class RIM Study." 2025. gens-associates.com/world-class-rim-research/.
5. International Council for Harmonisation. "M11 Template: Clinical Electronic Structured Harmonised Protocol (CeSHarP)." Step 4 final, November 19, 2025. database.ich.org/sites/default/files/ICH_Step4_M11_Final_Template_2025_1119.pdf. FDA: fda.gov/regulatory-information/search-fda-guidance-documents/m11-technical-specification-clinical-electronic-structured-harmonised-protocol.
6. U.S. Food and Drug Administration. "Knowledge-aided Assessment & Structured Application (KASA)." fda.gov/media/162865/download. PQ-CMC: fda.gov/media/161618/download.
7. European Medicines Agency. "Data on medicines (ISO IDMP standards): Overview." ema.europa.eu/en/human-regulatory-overview/research-development/data-medicines-iso-idmp-standards-overview. FDA: fda.gov/industry/fda-data-standards-advisory-board/identification-medicinal-products-idmp.
8. U.S. Food and Drug Administration. "Electronic Common Technical Document (eCTD) v4.0." fda.gov/drugs/electronic-regulatory-submission-and-review/electronic-common-technical-document-ectd-v40. RPS basis: "Specifications for eCTD v4.0 Validation Criteria." fda.gov/media/179723/download.
9. International Council for Harmonisation. "M4: The Common Technical Document." ich.org/page/ctd. Council for International Organizations of Medical Sciences. cioms.ch.
10. International Council for Harmonisation. "E2B(R3) Individual Case Safety Report (ICSR) Specification and Related Files." ich.org/page/e2br3-individual-case-safety-report-icsr-specification-and-related-files.
11. U.S. Food and Drug Administration. "eCTD Resources" (SPL, an HL7-approved document markup standard adopted by FDA). fda.gov/drugs/electronic-regulatory-submission-and-review/ectd-resources.
12. CDISC. "Unified Study Definitions Model (USDM)." cdisc.org/ddf. TransCelerate BioPharma. "Digital Data Flow." transceleratebiopharmainc.com/initiatives/digital-data-flow/.
13. International Council for Harmonisation. "Q12: Technical and Regulatory Considerations for Pharmaceutical Product Lifecycle Management." Step 4, November 2019. database.ich.org/sites/default/files/Q12_Guideline_Step4_2019_1119.pdf.
14. PMC (National Library of Medicine). "Structured content and data management: enhancing acceleration in drug development through efficiency in data exchange." ncbi.nlm.nih.gov/pmc/articles/PMC10164450/.
15. Docuvera. Reported deployment outcome, Clinical Overview update, governance-first AI pilot (company-reported; no public URL).
16. Astrix. "Structured Content Authoring (SCA) in Clinical and Regulatory Documents." astrixinc.com/blog/structured-content-authoring-sca-in-clinical-and-regulatory-documents/. Content Rules. "How to Tailor Structured Content Strategy in Pharma and Biotech." contentrules.com/tailoring-structured-content-strategy-in-pharma-and-biotech/.
17. Gravitare Health. "A Digital Health Information Journey" (IMI public-private partnership; the Gravitare Lens, G-lens®, focuses approved FHIR ePI). gravitatehealth.eu.