



ADLIB

A Practical Guide For Life Sciences

Defensible AI in Life Sciences

What Your Leadership Wants,
And What Your Team Needs to Know

Inside this guide

- Why AI keeps stalling in regulated industries (it is **not** the model)
- Why your trust concerns are valid, and how they can be resolved
- How Adlib differs from OCR, Data Extraction, and tools built into your systems
- Key use cases: eCTD, FDA correspondence, TMF, clinical AI readiness
- The architecture that makes AI deployable in regulated environments
- A self-assessment checklist your team can use this week

Why AI Initiatives Keep Stalling in Life Sciences

Leadership has seen the demos. The board has approved the budget. A pilot ran and looked promising. Then something happened between the demo and production, accuracy

dropped, hallucinations appeared, the validation team flagged concerns, and the audit trail didn't exist. Sound familiar?

THE BOTTLENECK IS NOT THE MODEL

Most enterprise AI failures in life sciences are not caused by choosing the wrong LLM. They are caused by **documents that were never prepared to be used as AI inputs**. Clinical study reports, eCTD submissions, FDA correspondence, batch records, and trial master file documents are structurally complex, multi-modal, and evidence-bearing. They were not designed to be flattened into a token stream.

When an AI system ingests a 600-page regulatory submission the same way it ingests a customer support FAQ, the chain of custody breaks before the model ever sees the content. Tables become noise. Footnotes vanish across page boundaries. Diagrams disappear. Metadata is dropped. Versions are lost.

THE 5 PHASES OF A STALLED AI PROGRAM

1 The Pilot Wins the Budget: Clean documents, contained use case, impressive accuracy. Budget approved.

2 The Production Wall: Real documents arrive. Every format, every scan quality, every exception the business has accumulated. Accuracy collapses.

3 The Validation Tax: Human review is added permanently. SMEs pulled into queues. ROI evaporates.

4 The Audit Reckoning: Compliance asks: How is output traced to source? Which version was used? How do we reproduce this in 18 months? No answers exist.

5 The Quiet Shelving: The use case becomes "experimentation." A new program is announced with a similar plan.

Why AI Initiatives Keep Stalling in Life Sciences

WHAT LEADERSHIP IS ACTUALLY ASKING FOR

When executives say “we need to look into AI,” they are asking three questions that no model benchmark answers: ***Can we trust this? Can we defend this? Will this scale?*** A program that cannot answer all three does not deploy, it demonstrates.

The good news is that the problem is architectural, not fundamental, and it can be solved. The solution sits in the layer beneath the model, the layer that prepares documents before they reach AI.

Key insight

*AI does not fail because models are incapable.
It fails because documents arrive unprepared.
The teams are talented. The models are capable. The documents are the problem.*

Your Trust Concerns Are Valid Here Is Why They Can Be Resolved

If you are hesitant about AI in regulated workflows, your instincts are correct. The concern is not paranoia, it is regulatory

awareness. But the solution exists, and it does not require waiting for a better model.

WHAT REGULATORS ACTUALLY REQUIRE

FDA, EMA, ICH, and the GxP framework all converge on the same five requirements for any document used in a regulated decision:

- ▶ **Preserved without alteration**
- ▶ **Provenance traceable to its origin**
- ▶ **Integrity verifiable**
- ▶ **Version unambiguous**
- ▶ **Decision derived from it reproducible on demand**

Most AI systems today satisfy *none of these*. They produce answers that are confident, fluent, and well-formatted with no trail back to the source, no version record, no reproducibility guarantee. That is the trust gap. And it is not a model limitation but an architecture limitation.

THE “PLAUSIBLE ANSWER” PROBLEM

The most dangerous failure mode in regulated AI is the hallucination that *looks right*. A model will summarize a document it never fully read, cite a table it never parsed, and explain

a diagram it never received, all in clean, confident prose. Reviewers without time to verify accept the answer. *Fluency is not accuracy*.

IN LIFE SCIENCES, DOCUMENTS ARE EVIDENCE, NOT JUST DATA

Data vs. Evidence

A pivot table of dosages is **data**. The clinical study report that supports the dosage decision is **evidence**.

Data is read for information. Evidence is read for proof, and proof must be preserved intact, attributable to a source, and reproducible by anyone who looks.

Most AI systems are built for data. Very few are built for evidence. That gap turns successful pilots into failed audits.

Your Trust Concerns Are Valid Here Is Why They Can Be Resolved

WHAT “DEFENSIBLE AI” ACTUALLY MEANS (OPERATIONALLY)

Defensible AI has a precise definition. Every AI output must be:

- ▶ **Traceable to a specific page, paragraph, table cell, or image**
- ▶ **In a specific version of a specific document**
- ▶ **Ingested through a documented channel, on a documented date**
- ▶ **Reproducible six months later by someone not on the original project**

If your AI cannot point to the page, you do not have an AI answer. You have an opinion. In a regulated enterprise, opinions do not survive an audit.

The path forward

Trust in regulated AI is not built at the model layer. It is built at the document layer, the system that normalizes, validates, traces, and governs documents before they reach any AI. Organizations that solve the document layer first are the ones that get AI into production.

The Root Cause: Your Documents Are Not AI-Ready

“AI-ready” is not the same as *digitized*. A PDF that exists in your document management system is not ready for AI. A scanned clinical study report is not ready for AI. Neither is

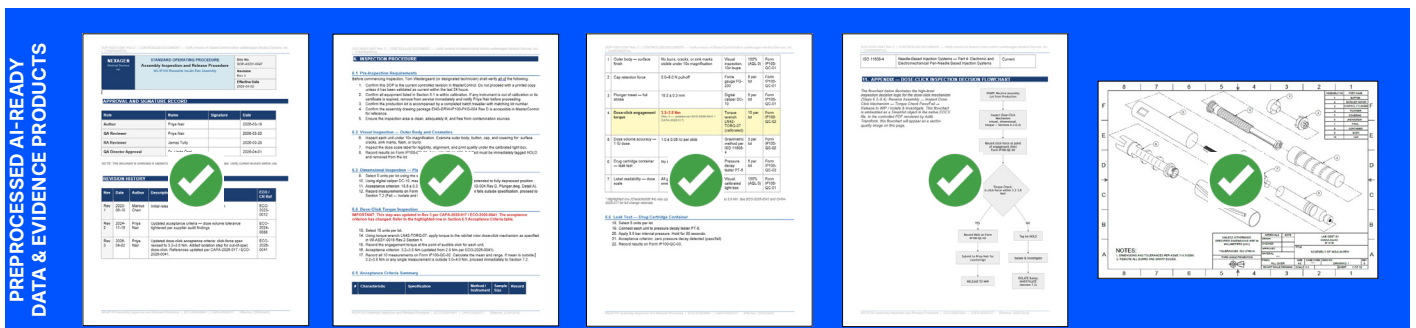
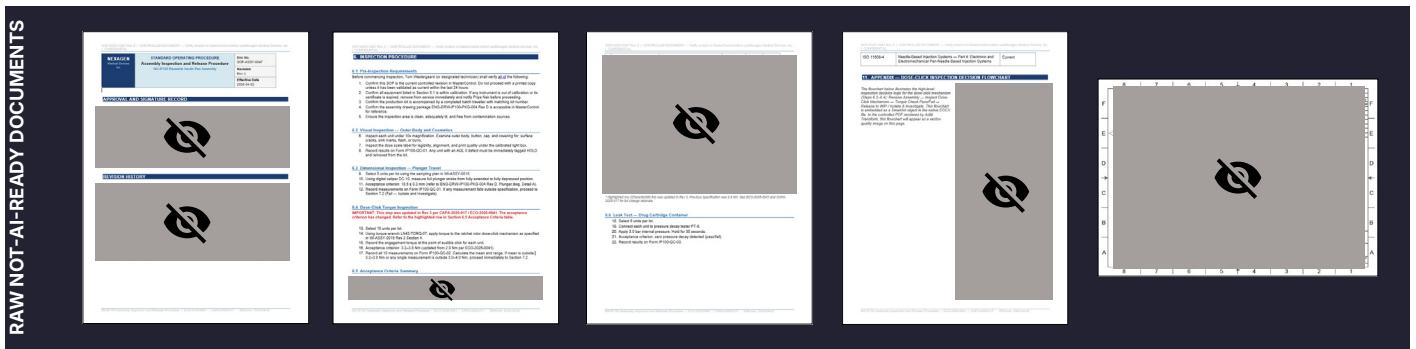
a document that has been converted with a basic PDF tool, extracted with OCR, or chunked into a RAG pipeline.

WHAT LLMS ACTUALLY SEE

A large language model does not see document structure. It sees a flattened token stream. A table becomes a smear of cells with no row context. A header becomes indistinguishable from body text. A diagram is simply absent, the model never receives it. A cross-reference in a 600-page clinical study

report pointing to Appendix F is broken the moment the document is chunked.

The model does not warn you when content is missing. It produces an answer from whatever it received, and the gap between that answer and the actual document is invisible until someone checks.



- ▶ All content is machine-readable: text, tables, handwritten annotations, CAD images, and diagrams
- ▶ Document structure, hierarchy, and cross-references are preserved so AI sees the whole, not fragments
- ▶ Every output is traceable to source, reproducible on demand, and ready for inspection

The Root Cause: Your Documents Are Not AI-Ready

WHY OCR IS NOT ENOUGH

OCR produces text. It does not produce trust.

OCR gives you characters on a page. It does not give you structure (headers vs. body), table integrity, cross-reference linkage, version provenance, metadata, confidence scoring, or any indication of what the model is missing. A document processed through OCR alone is searchable, it is not audit-ready, not AI-ready, and not defensible.

WHY CHUNKING AND RAG FAIL IN REGULATED CONTEXTS

RAG (Retrieval-Augmented Generation) has become the default workaround for large documents. It works well when documents are short, self-contained, and topically uniform. Clinical trial documents, regulatory submissions, and FDA correspondence are none of those things.

- Section 7 of a clinical study report assumes Section 3. Chunking breaks that relationship.
- Table 4.2 is the data behind the conclusion in Chapter 5. Different chunks, never reunited.

- Footnote 19 is the reservation that makes the headline number defensible. Lost across a page boundary.
- Appendices that justify a conclusion are filed in a different chunk and never consulted.

RAG is a useful tool inside a real document architecture. As the architecture, it is a liability. The question is not “is our RAG tuned?”, it is “what feeds our RAG?”

WHAT AI-READY ACTUALLY MEANS FOR A LIFE SCIENCES DOCUMENT

Requirement	What it means
Cleaned	Noise removed; scans de-skewed; distorted images fixed
Structured	Headers, sections, tables, figures, footnotes recognized and preserved
Extracted	Critical fields isolated into structured data contracts
Validated	Outputs checked against business and compliance rules
Traceable	Every output linked back to source page, version, and ingestion date
Governed	Redaction, retention, and access enforced before the model sees content

How Adlib Differs From OCR, Data Extraction Tools, and PDF Tools in Veeva

Your organization already uses PDF conversion tools, may have data extraction capabilities, and almost certainly has document tools built into Veeva Vault or other

TMF, CTMS, RIM and QMS systems. Here is the honest comparison: what each does well and where each falls short in regulated AI workflows.

Capability	OCR / PDF Conversion	Data Extraction Tools	Document Accuracy Layer
Convert files to text/PDF	✓ Yes	✓ Yes	✓ Yes, pixel-perfect, fidelity-preserving
Preserve document structure (headers, tables, hierarchy)	× Partial or lost	× Inconsistent	✓ Full structural preservation
Distributed OCR with quality thresholds	× None	× Limited	✓ Distributed; low-quality routed to remediation
Per-document confidence scoring (TrustScore)	× None	× Basic	✓ Hybrid AI + rule-based confidence scoring
Cross-document validation and completeness gates	× None	× None	✓ Field-level and packet-level completeness checks
eCTD / FDA technical standard compliance	× None	× None	✓ FDA/EMA eCTD standards built in
Exception routing to human-in-the-loop	× None	× Limited	✓ Configured by document class and confidence threshold
Audit trail with provenance	× None	× Minimal	✓ Full traceable provenance; inspection-ready
Veeva Vault / RIM / eTMF integration	× Manual	× Custom build	✓ Native connectors; structured data product output
Context preservation for RAG and AI agents	× Breaks structure	× Partial	✓ Citation anchors, stable chunks, structured contracts
Governance: redaction, retention, access control	× None	× Limited	✓ Enforced at ingestion, before the model sees content

How Adlib Differs From OCR, Data Extraction Tools, and PDF Tools in Veeva

WHAT ABOUT TOOLS BUILT INTO MY SYSTEM OF RECORD?

Life sciences-native systems of record like Veeva Vault, MasterControl, ArisGlobal, etc, have document management, version control, and some basic rendering and conversion capabilities built in. These are designed for storage, workflow routing, and access control, not for making documents AI-ready. Adlib is

not a replacement for Veeva or any other core system of record; it is the accuracy layer that sits upstream, ensuring that every document flowing into your systems (and into AI systems from there) is normalized, validated, and traceable before it matters.

HOW ADLIB CONNECTS TO YOUR EXISTING STACK

Adlib is designed to integrate, not replace. Whether your team works through pre-built connectors, direct API calls, or AI-native

protocols, Adlib fits into the architecture you already have.

Three ways to connect Adlib to your environment:

▶ Out-of-the-Box (OOTB) Connectors

Pre-built integrations for the platforms life sciences teams already use: Veeva Vault, OpenText, Microsoft SharePoint, ECM and RIM platforms, CTMS, eTMF systems, ERP, and data lakes. Normalized, governed document outputs flow directly into your system of record, no custom integration work required.

▶ REST API

For teams building custom automation pipelines, orchestration workflows, or proprietary platforms, Adlib exposes a full REST API. Submit documents, receive structured outputs (PDF/A, JSON data contracts, confidence scores, provenance metadata), and integrate the accuracy layer into any workflow, such as n8n, UiPath, AWS, Azure, or your own stack.

▶ MCP (Model Context Protocol)

Adlib supports MCP, the emerging open standard for connecting AI agents and LLMs to external tools and data sources. This means AI agents (Claude, ChatGPT, Copilot, and others) can call Adlib directly at inference time: submit a document, receive structured, validated, citation-anchored content, and use it as grounded evidence in their response. This is the architecture that makes agentic regulatory workflows safe to deploy.

Whether your team is routing documents through Veeva today, building an agentic pipeline tomorrow, or somewhere in between, Adlib connects where you are.

Key Life Sciences Use Cases Where Adlib Delivers Value

Each of the use cases below represents a workflow where document quality and traceability directly determine regulatory outcome, submission success, or inspection

readiness. Adlib delivers value by making the documents in each workflow AI-ready, governed, and audit-defensible.

Use Case	What Breaks Today	What Adlib Does	Expected Outcomes
eCTD / Regulatory Submissions	Inconsistent formats; manual hyperlinking; metadata gaps; 60% of rejections from technical errors	Automates eCTD readiness: ingest → classify → render → validate → audit → output to RIM/eTMF with FDA/EMA compliance	50% less prep time ~\$250K saved per resubmission
FDA Correspondence (CRLs, 483s, Deficiency Letters)	Fragmented intake across portals/email; attachments detached; wrong version actioned; no deadline tracking	Agentic intake captures all correspondence types; extracts structured action items; validates completeness; routes exceptions	Faster response drafting; fewer missed commitments; defensible audit trail
Clinical Trial Document Intake & TMF/eTMF	Mixed formats from sites/CROs; manual quality checks; weak provenance; TMF gaps discovered at inspection	Normalizes all clinical packages; validates completeness per artifact type; publishes governed records to CTMS/eTMF	Inspection-ready TMF; reduced manual QC; full chain of custody
GxP Audit Trails & Inspection Readiness	Batch records, SOPs, and quality documents processed inconsistently; audit trails reconstructed manually	Applies stamps, signatures, barcodes; enforces regulatory format compliance; routes to ECM/ERP/SharePoint with full provenance	Born-audit-ready outputs; reduced manual tracking; faster inspection response
Clinical AI & RAG Readiness	LLMs given raw PDFs produce inconsistent results; hallucinations rise; RAG retrieves wrong chunks; outputs non-reproducible	Provides structured, citation-anchored, confidence-scored document inputs so AI agents and RAG pipelines work on trusted content	Consistent, reproducible AI outputs; measurable confidence; grounded evidence for clinical copilots

Outcome data from the eCTD use case is sourced from Adlib customer results.

FDA citation: Manual hyperlinking and bookmarking accounts for 50%+ of technical reject triggers in recent ANDA submissions (FDA, 2023).

The Architecture That Makes AI Deployable in Regulated Environments

Most organizations are not choosing between build and buy. They are choosing between two versions of build, one labeled “innovation” and one labeled “integration”, and discovering

too late that neither solves the document problem. The right frame is different.

THE THREE LAYERS OF AN ENTERPRISE AI STACK

Layer	Description
Use case / application layer	Your workflows, domain logic, regulatory models, competitive IP. <i>Build this, it is what differentiates you.</i>
Orchestration / model layer	LLMs, RAG pipelines, agent frameworks. <i>Buy this, models are commoditizing every quarter.</i>
Document Accuracy & Trust Layer	The system that turns messy, multi-format enterprise documents into trusted, structured, traceable inputs. <i>Buy this, no enterprise should build it from scratch. This is Adlib.</i>

SEVEN CAPABILITIES PRODUCTION-READY DOCUMENT AI REQUIRES

- 1 Universal ingestion**
every format, every channel, no exceptions; 95% is not production
- 2 Structural understanding**
headers, tables, figures, hierarchy preserved and machine-readable
- 3 Context preservation**
the right content assembled for the model in the right order, with relationships intact
- 4 Validation and accuracy controls**
deterministic checks; confidence scores that are calibrated, not theatrical
- 5 Traceability and audit**
every output linked to source artifact, version, and page automatically
- 6 Scale and exception handling**
accuracy must hold at millions of documents, not just the demo corpus
- 7 Governance and policy enforcement**
redaction, retention, and access enforced before the model sees content

The Architecture That Makes AI Deployable in Regulated Environments

HOW TO GET THERE WITHOUT A RIP-AND-REPLACE

Three-phase transition

- ▶ **Phase 1:** Pick your highest-stakes AI use case. Route its document inputs through an accuracy layer. Measure: exception rate, validation cost, audit readiness.
- ▶ **Phase 2:** Migrate the next two or three use cases to the same layer. Document unit economics. Compare to per-program cost of the previous approach.
- ▶ **Phase 3:** Make the accuracy layer the default. New use cases start on it. The architecture becomes the standard, not the project.

The first use case pays for the foundation. Every subsequent use case runs on it at a fraction of the cost.

WHAT ADLIB TRANSFORM DELIVERS

Adlib Transform is the Document Accuracy and Trust Layer, *the AI production layer*, that makes your AI work, scale, and defend itself without replacing the LLMs, RAG systems, Veeva Vault, or automation platforms already in place.

Adlib does not compete with your model. It makes your model deployable. It does not compete with your ECM or RIM. It makes what goes in and comes out of those systems trustworthy.

FROM THE FIELD | CUSTOMER STORY

WHAT SCALE ACTUALLY LOOKS LIKE

A global medical products manufacturer runs more than one hundred thousand documents through Adlib every month (manufacturing batch records, quality records, SOPs, audit and inspection reports) at a single class of regulated production site, repeated across sites worldwide.

The work is not glamorous. Adlib renders the batch records, applies stamps, signatures, and barcodes, ensures regulatory format compliance, routes the output to ECM, ERP, and SharePoint systems, and prints batch packets to printers on the manufacturing floor. A custom in-house application orchestrates the workflow to satisfy regulatory Data Integrity guidance.

“The rendering engine allows us to render hundreds of thousands of documents. Without this capability, we would have to manually assign a tracking number to each batch document.”

Why this matters here. This is what a document trust layer already looks like in life sciences. The enterprises that get AI right will be the ones that build their AI on top of a layer already proven at this volume.

Take This Back to Your Team: Self-Assessment and Next Steps

Before your next internal conversation about AI readiness, run through these questions. Each “No” or “Unsure” is a gap in your

document accuracy layer, and a risk in any AI program built on top of it.

AI READINESS SELF-ASSESSMENT (5 QUESTIONS)

Question	Yes	No / Unsure
Can you trace any current AI output back to a specific page in a specific version of a specific source document?	<input type="checkbox"/>	<input type="checkbox"/>
Can you reproduce that AI output 12 months later with no member of the original project team?	<input type="checkbox"/>	<input type="checkbox"/>
Has compliance, legal, or risk signed off on how AI outputs will be validated and audited?	<input type="checkbox"/>	<input type="checkbox"/>
Do you have a redaction, retention, and access policy enforced before the model sees the content?	<input type="checkbox"/>	<input type="checkbox"/>
If a regulator asked tomorrow, could you describe the document accuracy layer your AI sits on?	<input type="checkbox"/>	<input type="checkbox"/>

THREE ACTIONS YOU CAN TAKE THIS WEEK

1 Audit your document layer.

Map the document types and formats feeding your top two or three AI use cases. Count the exceptions. The gaps you find are your real AI readiness roadmap.

2 Reframe the build vs. buy conversation

Stop debating which model to use. Start asking which layers to build and which to buy. You build the workflows that differentiate you. You buy the infrastructure that defends you.

3 Stress-test for defensibility.

Pick one current AI output. Try to trace it back to its source documents, page numbers, and version history. If you cannot, you have found the work.



ADLIB

Book an AI Readiness Review

Bring your messiest document workflow (a CRL response backlog, a 483 remediation package, an eCTD preparation cycle) and we will map what can be automated, where trust thresholds belong, and what your document accuracy layer should look like.

You will walk away with:

- A reference workflow: ingest → validate → extract → assemble → deliver
- Defined trust controls (thresholds + exception routing)
- A metrics plan (pass rate, exception rate, accuracy uplift)
- A 90-day action plan with your AI Readiness Score

[Request a Review →](#)



Trusted by



© 2026 Adlib Software. All rights reserved.

This guide and its contents are the intellectual property of Adlib and are protected by copyright law. No part of this guide may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.