AI agents inside the drug maker

What they can read, what they can change, and who signs

IN BRIEF

01 Drug makers have started to put AI to work on their documents. One uses it to draft clinical study reports, with a person reviewing and approving each draft. Others have told the press they use it to choose trial sites and to keep track of the thousands of pages a regulator asks for.

02 The regulators have started to answer. In April 2026 the United States Food and Drug Administration warned a small manufacturer that had used AI agents to write its specifications, procedures and master records, and had not checked them. In Europe, a draft annex to the good manufacturing practice guidelines says large language models “should not be used in critical GMP applications”. It is still a draft.

03 This paper is about the regulated work inside a drug maker, a biotech, a contract research organisation or a laboratory: the reading, checking and drafting done in manufacturing, quality assurance, the quality control laboratory, regulatory affairs, pharmacovigilance, clinical operations, medical information, validation and supply. It sets out what an agent can read, how it reaches a company’s systems, what it may change without a person, and what the rules already say. It ends with seven questions to put to any supplier, and with a place to start.

04 The position throughout is simple. An agent prepares. A named person decides, and signs. Regulated records are read, never written by an agent on its own. The agent runs on infrastructure the company controls.

IN THIS PAPER

01 Where drug makers are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 7
06 A worked example: one mid-sized drug maker 10
07 Seven questions before an agent touches a regulated record 11
08 Where to start 11

PART 01

Where drug makers are today

50%

of those polled at a clinical trials congress named trust and regulatory uncertainty as the biggest barriers to AI in trials

47%

of organisations in a supplier's survey, with 66 responses, reported no AI use in clinical quality management

One

warning letter from the FDA, sent on 2 April 2026, with a section on AI in pharmaceutical manufacturing

About 1,100

staff of a life insurer in Asia for whom the platform has run in production, on the insurer's own cluster

Drafting comes first. Novo Nordisk uses a generative AI platform to draft clinical study reports and other clinical documents. The published account says it gets documentation "directly into human hands for review and approval", and quotes the company: "In a highly regulated industry, we can't just throw our data and information into a large language model and hope for the best" [1].

Trials and submissions follow. Reuters reported in January 2026 that Novartis had used AI to choose the sites for a trial, that Genmab plans to use agents for the work that follows a trial, and that AstraZeneca, Roche and Pfizer described tracking thousands of pages of documents for regulators. The chief medical officer of Novartis put it this way: "AI becomes augmenting intelligence, not artificial intelligence" [2]. The article does not say how the output is checked.

Trust is the brake. At a clinical trials congress in London in May 2026, a poll found trust and regulatory uncertainty to be the biggest barriers to AI in trials, "cited by half of respondents (50%)"; the sample size is not given [3]. A survey run by a software supplier, with 66 responses, found that "47% of organizations reported no AI use in clinical quality management today", and concluded that adoption "is being blocked not by technology, but by confidence" [4].

The regulator has written once. On 2 April 2026 the FDA sent a warning letter to a small maker of homeopathic drug products in Michigan. One section is headed "Inappropriate Use of Artificial Intelligence in Pharmaceutical Manufacturing". The firm had told investigators it used AI agents: "Specifically, you used AI to create drug product specifications, procedures, and master production or control records." The letter goes on: "If you use AI as an aid in document creation, you must review the AI generated documents to ensure they were accurate and actually compliant with CGMP. Your failure to do so is a violation of 21 CFR 211.22(c)." And should the firm resume production, "any output or recommendations from an AI agent must be reviewed and cleared by an authorized human representative of your firm's QU" [5].

What the letter is, and is not. It concerns one small firm, and it also cites insanitary conditions and unapproved new drugs. It applies good manufacturing practice law that already existed and creates no new rule for AI. One law firm's reading: "To be clear, FDA does not prohibit the use of AI to assist with cGMP activities" [6]. The lesson is narrower: the quality unit answers for every document, whoever or whatever drafted it.

What the published accounts leave out. Only one of the accounts we opened says how a draft is checked. None describes a step where a named person approves each change an agent makes to a regulated record. Beyond the warning letter, we found no public incident at a drug maker caused by an AI agent, and we do not claim one.

What we can speak to. We name no drug maker here, because we have none to name. What we can say is that the platform runs inside regulated companies of other kinds. In production at a top-tier bank in Asia, air-gapped, inside the bank's own perimeter. It has run in production for about 1,100 staff of a life insurer in Asia, on the insurer's own cluster, inside its network.

PART 02

The work nobody has packaged

We listed the jobs done inside a drug maker, from the batch record to the trial master file, and looked for each in the template libraries of the agent and workflow platforms. The libraries cover intake, drafting and content review. They thin out wherever a named person signs something.

WHERE TEMPLATES EXIST	WHERE NONE WAS FOUND
Adverse reaction intake and case records	**A batch record checked against the master record before the quality reviewer decides**
Summaries of adverse event reports	**A deviation written up in the site's own form**
Review of promotional and medical content	**An out-of-specification result set against the investigation procedure**
Consistency checks across a submission	**A cleaning or line clearance record checked against the procedure**
First drafts of submission documents	**A trial master file checked against the expected document list**
Drafts of protocols and trial materials	**A requirements traceability check for a validation**
	A change request read for the documents and filings it touches

Three findings matter more than the list.

First, the template nearest to a regulator acts with no approval step. One adverse reaction workflow codes each event with a hosted model, writes the case record to a hosted spreadsheet, and at its expedited step "Notifies PV team, triggers regulatory submission for serious/unexpected cases". The page describes no step where a person reviews a case before the record is written or the alert goes out [7]. Among the templates we opened, only a sample for reviewing medical content makes a person part of the design, and that sample says it "is not intended as a production-ready solution" [8].

Second, the templates we opened are built from hosted services: a hosted model, a hosted spreadsheet, a hosted mail service. That settles where patient, batch and unpublished product data go before anyone has asked.

Third, where an agent runs is not the only question. One suite supplier says its agents use models "hosted on Amazon Bedrock". A validation supplier says that for its cloud customers "all processing occurs within their own private cloud subscription". A smaller platform says: "Run on your private cloud, on-premise data center, or fully air-gapped environment" [9]. Most of the supplier pages we opened say nothing about where their agents run. So private and on-premises offers exist. The harder questions are which tools an agent may call, as whom, with whose approval, and on whose model. One laboratory software supplier says its agents "can register molecules, create notebook entries, process results, update workflows, and more", and that they "run automatically as work happens"; the page does not say that a person approves each action [9].

The gaps sit where the reading is slow, the records are controlled, and the decision already belongs to a named person. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a drug maker that means three kinds of thing.

Documents.

Standard operating procedures, specifications, validated methods, the approved product information, the protocol, regulatory guidance. These are indexed once as knowledge sources in the Knowledge Engine, with each reader's permissions kept, so an agent answers with the document, version and step it relied on and says when the documents are silent.

Systems of record.

Quality events, controlled documents, laboratory results, submissions, safety cases, trial documents. Most are reached through an interface the company already runs. The company gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval. Change operations are left out or held for a person.

Data for analysis, and the published literature.

The year's batches, stability results and safety cases, held in the data warehouse or a reporting database and read only. Published papers and public registries, reached through web search or their publishers' interfaces.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

Where the company's interface checks the caller, the agent acts as the person using it and cannot claim another.

FOUR WAYS INTO A DRUG MAKER'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A batch record, a deviation report, a source document for an adverse event, a site visit report	Nothing is connected. The safest start
02	**A knowledge source**	Procedures, specifications, methods, approved product information, the protocol, guidance	Permissions are kept per reader
03	**An imported description of an interface**	Quality events, documents, laboratory results, submissions, safety cases, trial documents	You choose the operations. Changes wait for approval
04	**A server from the tool catalogue, run in your cluster**	The data warehouse, reporting databases, web search of the published literature	It runs inside; check what each tool does

PART 03 · CONTINUED

31
interface descriptions taken from their publishers; the importer refused none

2,719
operations the importer found in them

1,505
of those operations read; 1,214 are sent as changes

566
of the 953 operations in three laboratory notebook and inventory products change something

What is published. We took 31 interface descriptions from their publishers: laboratory notebook and inventory products, public trial and literature registries, protein and pathway databases, and the standards for exchanging clinical datasets and tracing a pack through the supply chain. The platform's importer refused none and found 2,719 operations [8].

What that exercise taught us. By method, 1,505 of the 2,719 operations read and 1,214 are sent as changes. The public registries are built for reading: every one of the public trial registry's 9 operations reads, and 58 of the protein database's 59. The laboratory products lean the other way. The three laboratory notebook and inventory products that publish a description carry 953 operations between them, and 566 of those change something: they create, edit or delete. Import one whole and an agent could create, edit and delete laboratory records. And 25 of the 31 files declare no sign-in scheme, so the file alone does not say how a caller is checked.

What is not published. No supplier of a quality, batch record, laboratory information, regulatory, safety or clinical data system publishes an open description. Where nothing is published, the company supplies its own description, or the agent reads from the data warehouse. Nothing in this paper is a tested integration with a named product.

Regulated records. The safe design gives an agent the batch, laboratory, quality, safety and trial records read only.

So choose the operations at import, start with the ones that read, and add nothing that changes a record until the approval step has been seen working.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named person decides.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the deviation report, the case narrative, the response to an authority's question, the validation protocol, with each statement tied to its source.
03	**It lists** what is missing, what is outside its limit, what departs from the procedure.
04	**It stops** before anything that changes a record or reaches a regulator, a health professional, a trial site or a colleague. A tool that changes something waits for approval.
05	**It decides nothing about a person.** It does not assess, rank or score anyone.

THE AGENT PREPARES	THE PERSON DECIDES
A list of missing entries and values outside their limits in a batch record	**The quality reviewer reviews the record and decides on every item**
A deviation report in the site's form, with similar past events	**The quality owner classifies it, decides the root cause and closes it**
The fields of an incoming adverse event report, and what is missing	**A safety officer decides validity, seriousness, expectedness and causality**
A list of what a submission lacks against the content plan	**The regulatory lead decides what goes in, and submits**
A site visit report from the clinical research associate's notes	**The clinical research associate checks and signs it**

Five lines are drawn tighter than the rest.

Signatures. No agent signs a record. Review, approval and release signatures are made by the person, in the system of record, under the company's own procedure.

Batches and deviations. No agent releases or certifies a batch, and none closes a deviation, an investigation or a corrective action. It takes no accept or reject decision.

Regulators. No agent submits to a regulator.

Adverse events. No agent decides whether a case is valid, serious or expected, or whether a medicine caused the event.

People. No agent decides about a patient, a trial participant or a member of staff. The investigator decides about every patient.

Validation is the company's decision. Katonic's only certificate is ISO 27001. Whether and how an agent is validated for its intended use is the company's decision under its own quality system; the platform supplies the records that decision needs.

PART 05

What the rules already ask

This part needs a lawyer's read, and a quality professional's, before it is relied on. It reports what the texts say; it is not advice.

Manufacturing and quality

Release belongs to the Qualified Person. The European Union's good manufacturing practice guidelines have an annex on computerised systems, Annex 11, in operation since 2011. Where a system records certification and batch release, "the system should allow only Qualified Persons to certify the release of the batches" [10]. No agent in this paper certifies or releases anything.

A revised Annex 11 is in draft. The Commission consulted on it from 7 July to 7 October 2025. The draft says "All users should have unique and personal accounts", that the audit trail "should unambiguously capture the user who made a change", and that when work is outsourced "the regulated user remains fully responsible" [11]. Whether an agent's call should appear under the person's account, under one of its own, or both is for the company's quality unit to settle.

AI has a draft annex of its own. The same consultation carried a new Annex 22 on artificial intelligence. As of November 2026 it is a draft, and the Commission's list of annexes in force does not include it. The draft covers models used "in critical applications with direct impact on patient safety, product quality or data integrity", and then sets generative AI aside: "the document does not apply to Generative AI and Large Language Models (LLM), and such models should not be used in critical GMP applications." For uses that are not critical it says "personnel with adequate qualification and training should always be responsible for ensuring that the outputs from such models are suitable for the intended use, i.e. a human-in-the-loop (HITL)" [12].

The final text is not known. The European Medicines Agency held a workshop on the annex on 30 June and 1 July 2026 and says it "is still considering the implications of the stakeholder consultation results" [13]. We do not predict what the final text will say. What follows for an agent today is plain enough: support work, with a named and trained person responsible for the output, and never an accept or reject decision. Every manufacturing and quality agent in this paper is written that way.

Electronic records and signatures

A controlled record shows who changed it. Under the United States rule on electronic records, a system must use "secure, computer-generated, time-stamped audit trails to independently record the date and time of operator entries and actions that create, modify, or delete electronic records", and "Record changes shall not obscure previously recorded information" [14].

A signature belongs to one person. The same rule says: "Each electronic signature shall be unique to one individual and shall not be reused by, or reassigned to, anyone else" [14]. An agent can prepare a record for review. The signature, and what it means, belong to a person.

PART 05 · CONTINUED

Submissions to a regulator

The FDA's guidance on AI is a draft, and it leaves drafting out. In January 2025 the FDA published draft guidance on AI used to support regulatory decisions about drugs. It is still marked "Draft Level 1 Guidance" and not for implementation [15]. It sets out a seven-step process for showing that a model can be trusted for its context of use. It also says what it does not cover: AI "used for operational efficiencies (e.g., internal workflows, resource allocation, drafting/writing a regulatory submission) that do not impact patient safety, drug quality, or the reliability of results" [15]. Being outside that draft is not permission to file. Submission stays the company's act.

Two regulators have agreed ten principles. In January 2026 the FDA and the European Medicines Agency published ten guiding principles of good AI practice in drug development. They are not binding. The first is "Human-centric by design", and another asks that AI technologies follow "a risk-based approach with proportionate validation, risk mitigation, and oversight based on the context of use and determined model risk" [16].

Safety and trials

FIELD	WHAT WE READ	WHAT IT ASKS
Pharmacovigilance	EU good pharmacovigilance practices, Module VI (Rev 2)	A valid case needs "at least one identifiable reporter, one single identifiable patient, at least one suspect adverse reaction, and at least one suspect medicinal product". Serious valid cases are submitted "in no case later than 15 calendar days after initial receipt of the information" [17]
	EMA reflection paper on AI in the medicinal product lifecycle, September 2024	Where AI is used in pharmacovigilance, "it remains the responsibility of the MAH to validate, monitor and document model performance and include AI/ML operations in the pharmacovigilance system" [18]. The MAH is the marketing authorisation holder
	CIOMS Working Group XIV report, 2025	"The extent and nature of human oversight for an AI solution should be risk-based, incorporating quality assurance principles" [19]
Clinical trials	ICH E6(R3) Good Clinical Practice, adopted 6 January 2025	"The approach to validation of computerised systems should be based on a risk assessment that considers the intended use of the system" [20]

PART 05 · CONTINUED

The AI Act and privacy

The dates have moved. Regulation (EU) 2026/1744 set the AI Act's high-risk duties to apply from "2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III" and from "2 August 2028" for those tied to Annex I [21].

No high-risk use names this work. We searched the Act for "pharmacovigilance", "good manufacturing" and "clinical trial" and found none of them. Two points still touch a drug maker. Employment is listed: systems used "to monitor and evaluate the performance and behaviour of persons" at work are high-risk [22]. The training agent in this paper drafts content and assesses no one. And AI that is a safety component of a medical device can fall under Annex I. Which side a given agent falls on is a question for counsel.

Research is set apart. The Act "does not apply to AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development" [22]. An agent used in regulated operations is not that.

People must be told. Systems "intended to interact directly with natural persons" must be built so that those persons "are informed that they are interacting with an AI system", unless that is obvious [22]. In this paper a person sends every reply to a health professional.

Health data is a special category. The GDPR prohibits the processing of "data concerning health" unless a listed exception applies [23]. Case reports and trial records are health data. And a person "shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her" [23].

We did not open the rules of the United Kingdom, Japan, China, India or the Gulf states, or the guidance of the World Health Organization, and we did not confirm the date from which ICH E6(R3) applies in the European Union. A company should ask its own counsel.

What the rules have in common

01	**A named person answers** for every record, release, submission and safety decision.
02	**Every change is traceable** to a person and a time, and a signature belongs to one individual.
03	**Oversight follows risk,** and the more a use bears on a patient or on product quality, the less an AI model may decide.
04	**The regulated company stays responsible,** whoever supplied the system.

PART 06

A worked example: one mid-sized drug maker

Take a made-up drug maker: about 1,500 people, two manufacturing sites, a dozen marketed products. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01	**On the shop floor,**	an operator asks what the site's procedure says about a gowning step. The agent answers from the site's own procedure and shows the document, version and step. Asked about something the procedures do not cover, it says so and tells the operator to ask the supervisor.
02	**In batch record review,**	the agent reads a completed record against the master record and lists three missing entries and one value outside its limit, each with its place. It runs beside the manual review, not in place of it. The quality reviewer still reviews the record as the procedure requires, and decides on every item.
03	**In quality assurance,**	a supervisor's first account of a deviation is turned into the site's report form, with the batches touched and two similar past events. The quality owner classifies it and decides the root cause.
04	**In the safety team,**	the agent screens the week's literature search and sets aside 12 articles that may hold a case, with the passage for each. A safety reviewer reads all 12 and samples the rest.
05	**In regulatory affairs,**	the agent checks a variation's documents against the content plan and lists two that are out of date. The regulatory lead decides what goes in, and submits.

Every one of these reads documents or records the company already has, changes no record by itself, signs nothing, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches a regulated record

01 **Where does it run, and whose model is it?**
On our infrastructure, in our private cloud, or in the supplier's service? Who serves the model? Can it run with no connection to the internet?

02 **What identity does it act under?**
Does it act as the person using it, with their permissions, or as one all-seeing account?

03 **What can it change without a person?**
Which tools change a quality event, a laboratory result, a safety case or a filing, and does each wait for approval?

04 **Who signs?**
Show us that no tool signs, releases, closes or submits, and that the signature is made by the person in the system of record.

05 **How is each call recorded?**
Who asked, which tool ran, what was decided, and for how long is that kept? Can we switch off one agent, one tool, or all of them, at once?

06 **How is a change controlled?**
When the model, the prompt or a tool changes, who approves it, and what record is there for our change control?

07 **What do you say in writing about validation?**
Which part is ours, which is yours, and what records do you supply for it?

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and no patient or batch waits on the agent.

Questions on your own procedures

Operators ask what the site's own procedures say and get the answer with the document and step cited. No system of record is touched, so it can run on the first day.

Batch record review

The agent lists every missing entry and value outside its limit. The quality reviewer decides.

Literature screening

The agent sets aside the articles that may hold a case. A safety reviewer reads them and samples the rest.

One product, one site, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/LIFE-SCIENCES

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 Anthropic customer story, "Novo Nordisk", undated (copyright 2026). https://claude.com/customers/novo-nordisk

02 Reuters, "Drugmakers turn to AI to speed trials, regulatory submissions", 26 January 2026, as carried by Yahoo Finance. https://finance.yahoo.com/news/drugmakers-turn-ai-speed-trials-110213746.html

03 Pistoia Alliance, "Half of clinical trial professionals cite trust and regulatory gaps as barrier to AI adoption, Pistoia Alliance poll finds", 28 May 2026. https://pistoiaalliance.org/news/half-of-clinical-trial-professionals-cite-trust-and-regulatory-gaps-as-barrier-to-ai-adoption-pistoia-alliance-poll-finds/

04 Montrium, "The State of AI in Clinical Quality", Industry Report 2026 (66 responses). https://www.montrium.com/the-state-of-ai-in-clinical-quality

05 US Food and Drug Administration, Warning Letter 320-26-58 to Purolea Cosmetics Lab, 2 April 2026. https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/purolea-cosmetics-lab-722591-04022026

06 DLA Piper, "Life Sciences and AI alert: FDA Warning Letter highlights risks of using AI in drug manufacturing", 21 April 2026. https://www.dlapiper.com/en/insights/publications/2026/04/fda-warning-letter-highlights-risks-of-using-ai-in-drug-manufacturing

07 n8n workflow template 10493, "Process ADR intake and PV cases with Claude AI and Google Sheets", undated. https://n8n.io/workflows/10493-process-adr-intake-and-pv-cases-with-claude-ai-and-google-sheets/

08 Template libraries of agent and workflow platforms, searched 22 November 2026, including the AWS sample "Medical Content Review", https://github.com/aws-samples/sample-medical-content-review-agent ; and 31 interface descriptions fetched from their publishers the same day and run through the platform's importer (Katonic research).

09 Supplier pages: Veeva, "Veeva AI Agents to Be Released Across All Veeva Applications", 14 October 2025, https://www.veeva.com/resources/veeva-ai-agents-to-be-released-across-all-veeva-applications/ ; ValGenesis, "Ensuring Trust in AI: The ValGenesis Approach to Validating AI Agents and Functions for Validation and Qualification in Life Sciences", 3 October 2025, https://www.valgenesis.com/editorial/ensuring-trust-in-ai-the-valgenesis-approach-to-validating-ai-agents-and-functions-for-validation-and-qualification-in-life-sciences ; ibl.ai, "AI Platform for Pharmaceutical & Biotech", last updated 12 March 2026, https://ibl.ai/resources/enterprise/pharmaceuticals.md ; Benchling, "Introducing Benchling Agents, your infinite program team", 27 August 2026, https://www.benchling.com/blog/introducing-benchling-agents-your-infinite-program-team

10 European Commission, EudraLex Volume 4, "Annex 11: Computerised Systems", revision January 2011. https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf

11 European Commission, draft revised Annex 11, stakeholders' consultation on EudraLex Volume 4 (Chapter 4, Annex 11 and new Annex 22), 7 July to 7 October 2025, sections 2.6, 11.1 and 12.2. https://health.ec.europa.eu/document/download/40231f18-e564-4043-94de-c031f813d38b_en?filename=mp_vol4_chap4_annex11_consultation_guideline_en.pdf ; consultation page: https://health.ec.europa.eu/consultations/stakeholders-consultation-eudralex-volume-4-good-manufacturing-practice-guidelines-chapter-4-annex_en

12 European Commission, draft "Annex 22: Artificial Intelligence", same consultation, section 1 (Scope). https://health.ec.europa.eu/document/download/5f38a92d-bb8e-4264-8898-ea076e926db6_en?filename=mp_vol4_chap4_annex22_consultation_guideline_en.pdf ; list of annexes in force: https://health.ec.europa.eu/medicinal-products/eudralex/eudralex-volume-4_en

13 European Medicines Agency, "Good manufacturing practice: Multistakeholder workshop on expert contributions to artificial intelligence guidance development (Annex 22)", 30 June and 1 July 2026. https://www.ema.europa.eu/en/events/good-manufacturing-practice-multistakeholder-workshop-expert-contributions-artificial-intelligence-guidance-development-annex-22

14 US Code of Federal Regulations, Title 21, Part 11, sections 11.10(e) and 11.100(a), as at 1 November 2026. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11

15 US Food and Drug Administration, "Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products", draft guidance, January 2025. https://www.fda.gov/media/184830/download ; status page: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

16 US Food and Drug Administration and European Medicines Agency, "Guiding Principles of Good AI Practice in Drug Development", January 2026. https://www.fda.gov/media/189581/download

17 European Medicines Agency, "Guideline on good pharmacovigilance practices (GVP), Module VI (Rev 2)", 28 July 2017. https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-good-pharmacovigilance-practices-gvp-module-vi-collection-management-submission-reports_en.pdf

18 European Medicines Agency, "Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle", EMA/CHMP/CVMP/83833/2023, 9 September 2024. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

19 CIOMS Working Group XIV, "Artificial intelligence in pharmacovigilance", Geneva 2025. https://doi.org/10.56759/cdob6397 (read in the copy at https://dspace.ut.ee/bitstreams/a3721f9e-eb9f-4ecc-9d68-3dc9e33cbbcc/download)

20 International Council for Harmonisation, "Guideline for Good Clinical Practice E6(R3)", adopted 6 January 2025, section 4.3.4. https://database.ich.org/sites/default/files/ICH_E6%28R3%29_Step4_FinalGuideline_2025_0106.pdf

21 Regulation (EU) 2026/1744, 8 July 2026, amended Article 113. https://publications.europa.eu/resource/celex/32026R1744

22 Regulation (EU) 2024/1689 (AI Act), Article 2(6); Article 6(1) and Annex I; Annex III point 4; Article 50(1). https://publications.europa.eu/resource/celex/32024R1689

23 Regulation (EU) 2016/679 (GDPR), Article 9(1) and Article 22(1). https://publications.europa.eu/resource/celex/32016R0679