

Governing Artificial Intelligence in GxP Operations: From First Enforcement to Durable Assurance



A cGMP Consulting White Paper on AI Credibility, Human Oversight, and the Transatlantic Divergence – September 24, 2026

Artificial Intelligence | Quality Systems | Data Integrity | Regulatory Strategy

An educational resource for executive and technical leaders in FDA-regulated life sciences.

Executive Summary

Artificial intelligence has moved out of pilot projects and into production across regulated life sciences, handling batch review, deviation triage, laboratory analysis, documentation, and a widening set of quality and manufacturing tasks. In April 2026 the FDA issued its first warning letter citing the misuse of artificial intelligence in current good manufacturing practice documentation, and it set a concrete precedent: the output of an AI system does not become a controlled record until an authorized human has reviewed it. The precedent now exists. In most organizations, the governance needed to satisfy it does not.

That gap, between how fast AI is being adopted and how mature the governance around it is, is now a citable regulatory exposure. Firms are dropping AI into workflows that generate or influence controlled records faster than they are building the intended-use definitions, human oversight controls, model validation, and vendor oversight those systems need. The decision to use AI has mostly been made. The decision to govern it, in a lot of places, has not.

The regulatory foundation for governing AI here is taking shape, but it does not read the same in every jurisdiction. In January 2025 the FDA issued a draft guidance proposing a risk-based credibility assessment framework for AI models that support regulatory decisions about drug and biological product safety, effectiveness, and quality, built around the concept of context of use and a structured, seven-step process. The finalized Computer Software Assurance guidance explicitly pulls AI and machine learning tools into scope when they are used in production or the quality system. Europe went the other way. The draft Annex 22 published in July 2025 would limit artificial intelligence in GMP-critical applications to static, deterministic models, keeping dynamic models, generative AI, and large language models out of such use. Operate globally and you now have to govern for two frameworks that agree on the principles and split on what is permitted at all.

This paper works through the regulatory landscape for AI in GxP operations, the things that make AI governance genuinely hard, the risks of ungoverned adoption, and the practices that build durable assurance. It gives you a roadmap for governance that keeps pace with adoption. Get credible, controlled AI governance in place and you can adopt AI safely and fast. Fail to, and you face either paralysis or the exposure the first enforcement action just made concrete.

Introduction

AI adoption in regulated life sciences has accelerated fast. These systems now help review batch records, triage deviations, analyze laboratory data, draft and summarize documentation, monitor processes, and support decisions across manufacturing and quality. The upside is real, faster cycle times, earlier detection of problems, relief from repetitive work. So is the risk, and it is a different kind of risk, because AI does not behave like the deterministic software regulated organizations have spent decades learning to govern.

Traditional computerized systems are predictable and they do not change after deployment. AI systems, machine learning ones especially, can produce outputs that are hard to predict, can carry the biases in their training data, and can shift their behavior as they learn or as their underlying models get updated. Those traits break the established methods of validation and control, and they raise questions traditional quality systems were never built to answer. When an AI system feeds a decision about product quality or patient safety, you have to be able to establish that the system is credible for that specific use and that real human oversight is in place.

The April 2026 warning letter made all of this concrete. By citing the misuse of AI in current good manufacturing practice documentation, the agency signaled that AI outputs face the same review-and-control expectations as any other data that becomes a controlled record. The principle is not new; it is longstanding requirements applied to a new technology. What is new is the certainty that the agency will enforce those requirements against AI, which ends any comfortable assumption that AI sits in a regulatory gray area.

This paper is for the quality, manufacturing, regulatory, and technology leaders deploying AI into regulated operations who have to keep adoption from outrunning governance. It focuses on the practical decisions that determine whether you can demonstrate credible, controlled use of AI when an inspector asks.

Industry Background

AI in regulated life sciences has grown from a niche research activity into a mainstream operational capability. The FDA has reported receiving hundreds of drug and biologic submissions with AI components since 2016, with an exponential increase over that stretch, and it has authorized a large and growing number of AI-enabled medical devices. That trajectory reflects both the maturing of the technology and the growing willingness of organizations to put it in consequential roles.

The AI in use varies enormously, from established, well-characterized static models doing narrow tasks to complex, adaptive systems whose behavior is much harder to predict. That range matters for governance, because the assurance that suits a simple, stable model is not the assurance a self-updating or barely-interpretable system needs, and because regulators are starting to draw lines along exactly that boundary. A one-size-fits-all approach to AI governance fails for the same reason one-size-fits-all software validation failed: effort has to be proportionate to risk.

The technical foundations of AI governance draw on disciplines the industry already knows, model validation, risk management, data integrity. A model's credibility depends on the quality and representativeness of its training data, whether the model actually fits its intended task, how rigorously it has been evaluated, and the controls over its use and its updates. None of that is entirely new. It extends familiar validation and data governance principles to a technology whose behavior is probabilistic instead of deterministic. The hard part is applying those principles in a way that accounts for what makes AI different, and doing it under frameworks that are still forming and do not yet fully agree with each other.

Regulatory Landscape

The regulatory framework for AI in GxP operations is assembled out of several documents, each covering a different piece of the problem. You have to understand how they fit together to build governance that satisfies the whole set.

The AI Credibility Framework for Regulatory Decision-Making

On January 6, 2025, the FDA issued a draft guidance titled Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products. It covers AI models used to produce information intended to support regulatory decisions about the safety, effectiveness, or quality of drugs and biologics, across nonclinical, clinical, post-marketing, and manufacturing phases. It does not cover AI used purely in discovery or in operational activities that do not affect patient safety, product quality, or study reliability.

The central concept is model credibility, meaning trust in a model's performance for a particular context of use. Context of use describes the specific role and scope of the model in addressing a question of interest. The guidance proposes a risk-based credibility assessment framework as a structured process: define the question of interest, define the context of use, assess model risk based on the model's influence and the consequence of the decision it supports, develop and execute a credibility assessment plan proportionate to that risk, and document the results. It lines AI governance up with established risk

management and model validation, and it establishes that the rigor of assurance should scale with the model's influence and the stakes of the decisions it informs.

Computer Software Assurance and AI in Production and Quality Systems

The Computer Software Assurance guidance, finalized September 24, 2025 and updated February 3, 2026 to align with the Quality Management System Regulation, explicitly states that its risk-based approach applies to artificial intelligence and machine learning tools when they are used as part of production or the quality management system. That brings AI in manufacturing and quality operations inside a familiar assurance framework, define intended use, assess the risk of failure, plan proportionate assurance activities, document the basis for confidence. For AI, that framework has to stretch to accommodate the dynamic nature of the technology, including the real possibility that a model's behavior changes over time.

The European Position: Draft Annex 22 and the Static Model Constraint

The European framework diverges from the American one in a way that carries direct strategic consequences. On July 7, 2025, the European Commission published a draft Annex 22 addressing artificial intelligence, alongside the draft revision of Annex 11 on computerised systems and a revised Chapter 4 on documentation. The consultation closed on October 7, 2025, and final publication is expected in mid-2026. Where the FDA's Computer Software Assurance guidance accommodates AI subject to risk-based assurance, the draft Annex 22 takes the harder line: it limits the use of AI in GMP-critical applications to static, deterministic models, and excludes dynamic models, generative AI, and large language models from critical use.

For a global operator, that is a big deal. An application that is permissible in the United States subject to appropriate assurance may fall outside the permitted envelope for GMP-critical use in Europe. This is not a difference in how much assurance you need. It is a difference in what is allowed at all. A firm that deploys a generative model into a critical quality workflow on the strength of a sound credibility assessment can satisfy the American framework while putting its European operations outside the draft's boundary. Because this is a deployment strategy question and not a control design question, it needs an explicit decision, not a technical fix, and you should track the final Annex 22 text closely, since the position it lands on will shape global AI strategy for anyone inspected in Europe.

Two things are worth underlining. The constraint applies to GMP-critical applications, so the line between critical and non-critical use becomes a governance decision with real teeth, which raises the stakes on precise intended-use and context-of-use definition. And the draft is not final, so plan against the direction rather than the exact wording, watching the final text while designing deployment strategy with the constraint in mind.

Artificial Intelligence in Drug Manufacturing and Device Software

The FDA addressed AI in drug manufacturing in earlier work, recognizing its role in advanced manufacturing and the questions it raises for process control and quality. On the device side, the agency has issued recommendations on predetermined change control plans for AI-enabled device software functions and on their lifecycle management, tackling the problem of systems designed to change after deployment. If you develop software as a medical device or combination products, those documents are

directly relevant, and their approach to managing change in AI systems informs governance more broadly, because assuring a system that updates is a problem common to every AI deployment.

Data Integrity, Records, and Human Oversight

The requirements that govern data integrity and controlled records apply to AI outputs. When an AI system generates or influences a record that satisfies a regulatory requirement, that record faces the same expectations of accuracy, attributability, and review as any other controlled record. The April 2026 warning letter applied this directly, establishing that AI output requires authorized human review before it becomes a controlled record. International frameworks reinforce the point, with the draft revision of Annex 11 expanding audit trail scope and access controls for computerised systems generally, while the draft Annex 22 addresses AI specifically. The ISPE GAMP Guide on Artificial Intelligence, published in July 2025, gives you an industry framework that complements both.

Two governing principles now apply, and they are not identical across jurisdictions. Everywhere, the rigor of AI assurance should be proportionate to the model's influence and the consequence of the decisions it supports, and human oversight is required wherever an AI system's output affects a quality decision or becomes a controlled record. In Europe, the draft Annex 22 adds a further constraint, limiting GMP-critical applications to static, deterministic models. Operate globally and you have to govern for both.

Current Industry Challenges

Governing AI in regulated operations throws up challenges that are distinct from traditional software, spanning the definition of intended use, the validation of learning systems, the design of human oversight, and the oversight of vendors.

Adoption Outpacing Governance

The most fundamental challenge is that adoption is running ahead of governance. Business units deploy AI tools to grab efficiency gains, sometimes before the quality organization has established the intended-use definitions, risk assessments, and controls those tools require. What you end up with is a growing population of AI systems in regulated workflows whose governance is immature or missing, which is exactly the condition the first enforcement action addressed. Closing that gap means bringing governance into the adoption process, not bolting it on afterward.

Defining Intended Use and Context of Use

Sound governance rests on a precise definition of what an AI system is meant to do and in what context. That is harder for AI than for deterministic software, because AI systems are often flexible and get applied to tasks beyond what they were assessed for. Without a clear context of use, you cannot assess model risk credibly, and you cannot set the boundaries within which the system's outputs can be trusted. Defining context of use, and then enforcing it, is a central governance discipline.

Validating Systems That Learn and Update

Traditional validation assumes a system, once validated, does not change. AI systems may change as they learn or as their underlying models get updated, which blows up that assumption. Establishing confidence in a system whose behavior can evolve takes a lifecycle approach to assurance, controls over how and when models change, and mechanisms to detect performance drift. Plenty of quality systems have no process to manage that kind of change, and it gets worse when the model comes from a third party whose update schedule you do not control.

Designing Effective Human Oversight

Human review of AI outputs is clear in principle and demanding in practice. Effective oversight means reviewers understand the system well enough to actually evaluate its outputs, review is genuine rather than a rubber stamp, and the workflow makes review practical at the scale and speed the AI runs at. There is a specific trap here: AI, by serving up confident, plausible outputs, can quietly erode the diligence of the human reviewer, and governance has to actively push back against that. Designing oversight that is both effective and sustainable is a real challenge, not a checkbox.

Divergent Expectations Across Jurisdictions

For firms inspected by multiple authorities, AI governance has to satisfy frameworks that do not fully agree. The American approach permits AI in production and quality systems subject to risk-based assurance, while the draft European Annex 22 would limit GMP-critical applications to static, deterministic models and shut dynamic and generative systems out of such use. Govern to the American framework alone and you can deploy a system that is well assured, well overseen, and still outside the European boundary. That is a strategic conflict, not a control gap, and you cannot fix it by strengthening assurance, because the constraint is on the category of model, not the quality of its validation.

Overseeing Third-Party and Embedded AI

A lot of the AI entering regulated operations is embedded in vendor systems or delivered as a service, which means you did not build the model, cannot fully inspect it, will not see its training data, and do not control its updates. Overseeing that takes contractual controls, supplier assessment, and ongoing monitoring, plus real clarity about who is responsible for what between you and the vendor. Traditional vendor qualification, built for stable systems, fits opaque, updating AI badly, and you have to adapt the oversight accordingly.

Risks of Poor Implementation

Ungoverned or poorly governed AI creates exposure across regulatory, quality, and business lines. The first enforcement action made the regulatory risk concrete, but the broader risks matter just as much.

Risk Category	Potential Consequence
Regulatory action	Warning letters and 483 observations citing inadequate control of AI outputs, following the precedent set in 2026, and the extended enforcement cycle that narrow responses now trigger.
Controlled record integrity	AI-generated content entering controlled records without authorized human review, compromising data integrity and calling into question the reliability of decisions that relied on it.
Erroneous decisions	Reliance on AI outputs that are inaccurate, biased, or applied outside their validated context of use, leading to incorrect quality or manufacturing decisions.
Model drift	Degradation of model performance over time or after updates, undetected because monitoring was inadequate, resulting in unreliable outputs.
Vendor exposure	Dependence on third-party AI whose behavior changes or fails outside the organization's control, with insufficient oversight to detect or manage the change.
Loss of confidence	Erosion of internal and regulatory confidence in the organization's ability to control its systems, extending scrutiny beyond the AI itself to the broader quality system.

AI carries a distinctive risk that deserves its own paragraph: false confidence. Because these systems produce fluent, plausible, internally consistent outputs, they can induce too much trust and dull the due diligence a human reviewer would normally apply. That makes inadequate human oversight especially dangerous, because the very quality that makes AI useful is the quality that undermines the review meant to control it. Governance has to treat human oversight not as a formality but as an active control that reviewers are trained and equipped to exercise for real.

As everywhere else in current enforcement, how you respond to a finding matters. Answer an AI-related observation with a narrow fix, instead of showing you have examined and strengthened your AI governance systemically, and you invite the inadequate-response letters that stretch the enforcement cycle. The expectation is evidence of a governed approach, not a point correction.

Best Practices

Durable AI governance rests on a handful of practices: bring governance into the adoption process, establish credibility proportionate to risk, and keep human oversight genuine across the whole lifecycle.

Establish an AI Governance Framework and Inventory

Stand up a governance framework that defines how AI gets evaluated, approved, controlled, and monitored, and keep an inventory of the AI systems in regulated use. The inventory should capture each system's intended use, context of use, risk classification, and oversight arrangements. That foundation makes governance visible and manageable, and it keeps AI systems from slipping into regulated workflows without review. Integrate governance into the adoption process, so new AI uses get assessed before they deploy rather than discovered after.

Define Intended Use and Context of Use Precisely

For each AI system, define precisely what it is meant to do and the context it operates in, following the context-of-use concept in the FDA's credibility framework. That definition sets the boundaries within which the outputs can be trusted and gives you the basis for risk assessment. Then build controls that stop the system from being applied outside its defined context, because use beyond the assessed context voids the assurance you did.

Assess Model Risk and Establish Credibility

Risk assessment should weigh the model's influence on a decision and the consequence of that decision, consistent with the credibility framework. High-influence, high-consequence uses earn rigorous credibility assessment, including evaluation of training data quality and representativeness, model performance, and the potential for bias. Lower-risk uses earn proportionate assurance. Document the credibility assessment so the basis for trusting the model is clear and defensible.

Design Genuine Human Oversight

Design human oversight as an active control. Define where human review is mandatory, make sure reviewers understand the system well enough to evaluate its outputs, and structure workflows so review is practical and genuine. Where an AI output influences a quality decision or becomes a controlled record, require authorized human review before the output takes effect, consistent with the enforcement precedent. And actively counteract the drift toward excessive trust, for instance by making sure reviewers have the information they need to challenge the AI's output rather than defer to it.

Manage the AI Lifecycle and Model Change

Because AI systems may change, adopt a lifecycle approach with controls over model updates, monitoring for performance drift, and re-assessment when a model changes materially. For systems designed to update, a predetermined change control approach, drawing on the FDA's work in this area, can define acceptable changes in advance and the conditions that force re-assessment. Monitoring has to be able to catch a model whose performance degrades, so you find the problem before it reaches a quality decision.

Decide the Global AI Position Deliberately


If you operate across jurisdictions, make an explicit decision about where AI may be deployed in GMP-critical applications, rather than letting deployment proceed under the more permissive framework and discovering the conflict at a European inspection. Because the draft Annex 22 constrains the category of model permitted in critical use, not the assurance applied to it, this belongs to strategy, not validation. Practical options: reserve dynamic and generative systems for non-critical applications; keep static, deterministic models for critical workflows across all sites; or accept differentiated deployment by jurisdiction with the governance complexity that brings. Each has costs. Make the choice deliberately and revisit it when the final Annex 22 text publishes. Under any of these, the boundary between critical and non-critical use becomes consequential, which again is why precise context-of-use definition matters so much.

Strengthen Vendor Oversight for AI Systems

For third-party and embedded AI, build oversight suited to opaque, updating systems. That means contractual provisions covering model changes, performance, and notification; assessment of the vendor's development and quality practices; and ongoing monitoring of the system's performance in your context. Define the division of responsibility between you and the vendor clearly, so you keep accountability for the outputs that enter your regulated processes, because the regulator will hold you to it regardless of who built the model.

Implementation Roadmap

The phased roadmap below lets you build AI governance that keeps pace with adoption and satisfies regulatory expectations.



Step 1: Assess the Current State

Inventory the AI already running across regulated operations, including embedded and third-party systems that were never formally recognized as AI. For each, capture intended use, context of use, and current controls. This reveals the real scope of the challenge and points you at the systems carrying the most exposure.

Step 2: Perform a Gap Assessment

Compare current governance against regulatory expectations, including the credibility framework, the Computer Software Assurance approach, and the requirements for controlled records and human oversight. Find the gaps in intended-use definition, risk assessment, human review, lifecycle management, and vendor oversight, and rank them by exposure.

Step 3: Develop the Governance Framework

Establish a governance framework that defines how AI is evaluated, approved, controlled, and monitored, and integrate it into the adoption process. Define roles and responsibilities, decision rights, and the documentation that will demonstrate governance. Make sure the framework scales assurance to risk so effort stays proportionate.

Step 4: Update Procedures

Revise procedures to require intended-use and context-of-use definition, model risk assessment, credibility assessment, human oversight, and lifecycle management for AI systems. Procedures should make clear when authorized human review is required and how model change is controlled.

Step 5: Train Personnel

Train the people who deploy, use, and review AI systems, building both the technical understanding to evaluate outputs and the awareness of the risks, including the pull toward excessive trust. Reviewers especially need to be equipped to exercise genuine oversight, and system owners need to understand their governance responsibilities.

Step 6: Assure and Deploy

For each AI system, complete the intended-use definition, risk assessment, and credibility assessment proportionate to risk, and put the human oversight and lifecycle controls in place before it enters regulated use. For systems already running, apply this process to bring them into compliance, starting with the highest-exposure ones.

Step 7: Monitor Performance and Drift

Establish monitoring that catches performance degradation and model drift, and that confirms human oversight is actually being exercised as designed. Monitoring should give leadership visibility into the health and control of the organization's AI systems.

Step 8: Govern Change and Improve

Manage model updates and changes through defined controls, re-assessing credibility when a model changes materially. Feed monitoring results, internal findings, and evolving regulatory guidance back into the framework, treating AI governance as an ongoing capability that matures as the technology and the regulatory environment develop.

Technology Considerations

The technologies relevant to AI governance include the AI systems themselves and the tools that support their governance.

AI and machine learning systems vary widely, and governance has to account for the differences. Well-characterized static models doing narrow, stable tasks can often be assured with proportionate effort, and they are the category the draft Annex 22 would permit in GMP-critical applications. Dynamic, adaptive, and generative systems need more rigorous credibility assessment and lifecycle control, and under the draft European position they would be excluded from critical use entirely. Systems with limited interpretability are the hardest case for human oversight, because reviewers may struggle to evaluate outputs they cannot readily explain, which puts a premium on clear context of use and on monitoring.

The infrastructure AI runs on, frequently cloud-based, brings the shared-responsibility considerations discussed under Computer Software Assurance. You have to understand which parts of the environment you control and which the provider controls, and oversee the provider accordingly. Integration between AI systems and the electronic quality and manufacturing systems that hold controlled records needs attention too, so the human review required before an AI output becomes a controlled record is enforced by the workflow rather than left to whoever happens to be at the keyboard.

Tools that support governance, systems for holding the AI inventory, managing risk and credibility assessments, and monitoring performance, can help you govern AI at scale. Monitoring tools that detect performance drift are especially valuable for systems that update, and workflow controls that enforce human review at the right points help make oversight consistent instead of discretionary. And like any computerized system, these governance tools should themselves be appropriately assured for their intended use.

Future Outlook

The regulatory framework for AI in life sciences will keep developing, and the picture is partial convergence with one significant open divergence. The FDA's credibility framework is expected to be finalized and may expand in scope, and industry frameworks, including the ISPE GAMP Guide on Artificial Intelligence, will keep maturing. On the principles, context of use, proportionate assurance, human oversight, the frameworks broadly agree. On what AI may be used for in GMP-critical applications, right now they do not. The final Annex 22 text, expected in mid-2026, is the most consequential near-term development for anyone inspected in Europe, and whether the static, deterministic model constraint survives consultation will shape global AI strategy directly.

Enforcement will follow adoption. The first warning letter addressing AI in manufacturing documentation is unlikely to be the last, and the agency's demonstrated willingness to enforce existing requirements against AI means you cannot lean on regulatory ambiguity. As AI takes on more consequential roles, credibility, human oversight, and control over change move further to the center of regulatory attention.

The organizations that thrive will treat AI governance not as a brake on innovation but as the thing that lets them adopt safely and fast. Credible, controlled governance is what lets you deploy AI with

confidence, capturing the benefits while managing the risks. Build that capability deliberately and you adopt AI faster and more safely than firms that either restrict it out of caution or deploy it without control, and you hold that advantage as AI becomes central to regulated operations.

How cGMP Consulting Can Help

Governing AI in regulated operations takes regulatory understanding, technical depth in AI and computerized systems, and practical experience in quality systems and data integrity. cGMP Consulting helps organizations build governance that keeps pace with adoption and satisfies the expectations recent enforcement has made concrete.

Our consultants help you inventory your AI systems, define intended use and context of use, and assess model risk and credibility in proportion to the influence and consequence of each system. We help design human oversight that is genuine and sustainable, establish lifecycle controls for systems that update, and develop the vendor oversight that opaque, third-party AI requires. We align governance with the FDA's credibility framework, the Computer Software Assurance approach, and the requirements for controlled records, and for firms inspected in Europe we advise on the strategic implications of the draft Annex 22 constraint on GMP-critical applications, helping determine where dynamic and generative systems can be deployed and where static, deterministic models are required. We help you prepare to demonstrate credible, controlled AI use during inspection, and for those responding to findings, we help build the systemic governance improvements regulators now expect.

Our approach is consultative and grounded in how AI use actually gets evaluated under current expectations. We aim to leave you with a governance capability you can sustain and extend, so you can adopt AI safely and quickly rather than either restricting it or deploying it without control.

Conclusion

AI is an operational reality in regulated life sciences now, and the first enforcement action addressing its use in manufacturing documentation has made the expectations concrete. An AI system's output faces the same review-and-control requirements as any other controlled record, and you have to be able to demonstrate that your use of AI is credible and controlled.

The path to durable AI governance is well defined. Bring governance into the adoption process. Define intended use and context of use precisely. Assess model risk and establish credibility in proportion to influence and consequence. Design genuine human oversight. Manage the lifecycle and control model change. Oversee the vendors that supply AI systems. And if you operate globally, make one more decision: where dynamic and generative systems may be deployed, given that the draft Annex 22 would confine GMP-critical applications to static, deterministic models. That is a strategic choice, not a validation exercise, and it should be made deliberately rather than discovered at inspection.

The organizations that succeed will treat AI governance as the enabler of confident adoption. The gap between the speed of adoption and the maturity of governance is now a citable exposure, but it is a

closable one. Close it deliberately and you are positioned to capture the benefits of AI while managing its risks, and you hold a durable advantage as AI becomes central to regulated operations.

If your organization is deploying artificial intelligence into regulated operations or seeking to ensure that its AI governance meets current regulatory expectations, the experienced consultants at cGMP Consulting can help assess your current state, identify governance gaps, and develop practical solutions tailored to your business objectives. To start a confidential conversation, complete the Contact Us form at cgmppconsulting.com/contact-us.

References

- European Commission. (2025). EudraLex Volume 4: Annex 22 Artificial intelligence (Draft, published 7 July 2025; consultation closed 7 October 2025). European Commission.
- European Commission. (2025). EudraLex Volume 4: Annex 11 Computerised systems (Draft revision, published 7 July 2025). European Commission.
- European Commission. (2025). EudraLex Volume 4: Chapter 4 Documentation (Draft revision, published 7 July 2025). European Commission.
- International Society for Pharmaceutical Engineering. (2025). GAMP guide: Artificial intelligence. ISPE.
- International Society for Pharmaceutical Engineering. (2022). GAMP 5: A risk-based approach to compliant GxP computerized systems (2nd ed.). ISPE.
- U.S. Food and Drug Administration. (2023). Artificial intelligence in drug manufacturing (Discussion paper). U.S. Department of Health and Human Services.
- U.S. Food and Drug Administration. (2024). Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions: Guidance for industry and FDA staff. U.S. Department of Health and Human Services.
- U.S. Food and Drug Administration. (2025). Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Draft guidance for industry (Docket No. FDA-2024-D-4689). U.S. Department of Health and Human Services.
- U.S. Food and Drug Administration. (2025). Artificial intelligence-enabled device software functions: Lifecycle management and marketing submission recommendations: Draft guidance for industry and FDA staff. U.S. Department of Health and Human Services.
- U.S. Food and Drug Administration. (2025). Artificial intelligence and medical products: How CBER, CDER, CDRH, and OCP are working together. U.S. Department of Health and Human Services.
- U.S. Food and Drug Administration. (2026). Computer software assurance for production and quality system software: Guidance for industry and FDA staff (Updated February 3, 2026). U.S. Department of Health and Human Services.
- U.S. Food and Drug Administration. (2018). Data integrity and compliance with drug CGMP: Questions and answers; Guidance for industry. U.S. Department of Health and Human Services.
- Federal Register. (2025, January 7). Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products; Draft guidance for industry; Availability; Comment request. U.S. Department of Health and Human Services.

- Niazi, S. K. (2026). A critical review of the FDA's draft guidance on artificial intelligence in drug and biological product regulation. *Journal of Chemistry*, 2026. <https://doi.org/10.1155/joch/5202999>
- International Council for Harmonisation. (2023). Q9(R1) Quality risk management. ICH.
- U.S. Government Publishing Office. (n.d.). Title 21 Code of Federal Regulations Part 11: Electronic records; Electronic signatures. *Electronic Code of Federal Regulations*.
- World Health Organization. (2016). Guidance on good data and records management practices (WHO Technical Report Series No. 996, Annex 5). World Health Organization.