EU GMP Annex 22: AI Compliance in Pharmaceutical Manufacturing

Published August 5, 2026 41 min read

EU GMP Annex 22: AI Compliance in Pharmaceutical Manufacturing

Executive Summary

EU GMP **Annex 22** is a proposed new annex to EudraLex Volume 4, the European Union's compiled Good Manufacturing Practice (GMP) guidance, describing proposed expectations for the use of **artificial intelligence (AI)** in GMP-regulated pharmaceutical manufacturing. It was drafted jointly by the European Medicines Agency (EMA) and Pharmaceutical Inspection Co-operation Scheme (PIC/S) GMP/GDP Inspectors Working Group, alongside a heavily revised **Annex 11** (Computerised Systems) and Chapter 4 (Documentation), and released for a three-month public consultation running from 7 July 2025 to 7 October 2025 (Source: picscheme.org). As of the August 2026 publication date of this report, the consultation is formally closed but no final, adopted guidance has replaced the draft; the binding GMP obligations arise from the applicable EU legislation, while EudraLex Volume 4 provides guidance for interpreting GMP principles and guidelines (Source: health.ec.europa.eu); EMA convened a further multistakeholder expert workshop on 30 June and 1 July 2026 specifically to help shape a risk-based approach for generative AI ahead of finalization (Source: ema.europa.eu).

The draft annex, at roughly six pages, is scoped narrowly and deliberately (Source: bioprocessonline.com). It applies only to **static ("frozen") AI/machine learning models with deterministic outputs** used in critical GMP applications directly affecting patient safety, product quality, or data integrity, explicitly excluding **dynamic (continuously self-learning) models**, **probabilistic-output models**, **generative AI**, and **large language models (LLMs)** from critical use (Source: picscheme.org). The draft sets out technical controls that go beyond generic Annex 11 computerised-systems validation: **explainability** techniques such as feature attribution (SHAP values or LIME) during model testing (Source: health.ec.europa.eu); confidence-score logging for each prediction or classification **where applicable**, and consideration of an "undecided" outcome when a score is very low (Source: health.ec.europa.eu); pre-deployment change control (Source: health.ec.europa.eu); and ongoing monitoring for input data drift (Source: health.ec.europa.eu). PIC/S summarizes the annex as foreseeing "continuous oversight of AI systems, including change control" and model performance monitoring (Source: picscheme.org).

The regulatory momentum extends well beyond the EU. In April 2026, the US Food and Drug Administration (FDA) issued what appears to be its first cGMP warning letter explicitly citing "overreliance on artificial intelligence" in drug manufacturing operations, to a firm that had used AI agents to generate specifications, procedures, and master production records without adequate human review (Source: fda.gov) (Source: raps.org). ICH published a Reflection Paper on 8 October 2025 naming AI-based process modelling as a target for future harmonized manufacturing guidance (Source: admin.ich.org), while WHO concluded in 2023 it was still too early to draft its own guideline (Source: cdn.who.int). Adoption on the ground is already substantial: a 2025 Rockwell Automation survey of 143 life-sciences manufacturing leaders found 95% are using or evaluating smart manufacturing technology including AI (Source: rockwellautomation.com), and named deployments at **Novo Nordisk**, **Merck (MSD)**, **Novartis**, **Roche/Genentech**, and **GSK** already apply machine learning to visual inspection, predictive maintenance, and manufacturing scheduling, with Merck's HawkAVI platform reporting a 50% cut in false product rejects (Source: aws.amazon.com).

Market forecasts for AI in pharmaceutical manufacturing vary considerably by methodology and scope: Fortune Business Insights projects the AI-in-pharma-manufacturing market to grow from $1.64 billion in 2026 to $12.90 billion by 2034 (Source: fortunebusinessinsights.com), while BCC Research forecasts the broader AI-in-pharmaceuticals market at $3.8 billion in 2025 rising to $15.2 billion by 2030 (Source: bccresearch.com). This report explains what Annex 22 covers, how it interacts with Annex 11, its explainability and confidence-score requirements, the static-versus-dynamic model distinction, the global regulatory landscape (PIC/S, FDA, ICH, WHO, MHRA, TGA), and a practical implementation checklist for manufacturers preparing for eventual finalization.

Introduction and Background

For more than a decade, EU GMP **Annex 11**, first issued in 2011, has governed "computerised systems" used in regulated pharmaceutical manufacturing and testing (Source: bioprocessonline.com). That annex was written for deterministic, rules-based software, not specifically for machine-learning models that learn patterns from data. Under draft Annex 22, only models with deterministic outputs are in scope for critical GMP applications. As pharmaceutical manufacturers increasingly deploy machine learning for visual inspection, predictive maintenance, batch release

support, and deviation triage, EU and international regulators concluded that Annex 11 alone did not adequately address the unique risks AI introduces: opaque decision logic, performance that can drift as input data changes, and the possibility that a model's confidence in its own output may not match its actual reliability.

In response, the European Commission's Directorate-General for Health and Food Safety, working with the EMA/PIC/S GMP/GDP Inspectors Working Group, released a package of three linked draft documents for consultation on 7 July 2025: a substantially revised Chapter 4 (Documentation), a substantially revised Annex 11 (Computerised Systems), and an entirely **new Annex 22** dedicated to artificial intelligence, jointly "drafted by the Inspectors' Working Group of the European Medicines Agency (EMA) and PIC/S"(Source: picscheme.org). The draft document itself confirms its novel status: under "reasons for changes," the drafting record simply states "not applicable (new annex)" (Source: health.ec.europa.eu). Annex 22 sits inside **EudraLex Volume 4**, the European Commission's compiled body of GMP guidance, which the Commission describes as containing "references to the applicable legislation and guidance for the interpretation" of GMP principles across the EU (Source: health.ec.europa.eu). Notably, Volume 4 still applies to veterinary medicinal products only until 15 July 2026, after which veterinary GMP guidance moves to a separate framework, a reminder that EudraLex itself is in a period of broader structural change alongside the Annex 22 work (Source: health.ec.europa.eu).

The three-month consultation period ran from 7 July 2025 to 7 October 2025, according to PIC/S's own announcement (Source: picscheme.org) and confirmed by the European Commission's official consultation record (Source: health.ec.europa.eu). As of this report's 5 August 2026 publication date, that same Commission page states plainly that "the response period for this consultation has ended," and no finalized replacement guidance has yet been published in its place. The legal GMP requirements derive from the applicable EU legislation; EudraLex Volume 4 provides guidance for interpreting those requirements (Source: health.ec.europa.eu). Trade press covering the initial release characterized the draft as compact and deliberately scoped: at roughly six pages, it was described as "the first significant foray into the AI" governance space by EU regulators (Source: bioprocessonline.com). The GMP-compliance trade publisher ECA Academy similarly noted the draft "has remained pleasingly short" relative to the much-expanded Annex 11 and Chapter 4 drafts released alongside it (Source: gmp-compliance.org).

Since the consultation closed, work has continued. EMA convened a two-day multistakeholder expert workshop on 30 June and 1 July 2026 specifically intended to "help shape a risk-based approach to the use of generative artificial intelligence" in medicines manufacturing, a technology category the July 2025 draft had explicitly excluded (Source: ema.europa.eu). This report is written for pharmaceutical manufacturing, quality assurance, and regulatory affairs professionals who need to understand what the draft Annex 22 actually requires, how it changes the relationship between AI systems and Annex 11 computerised-systems validation, and how to begin preparing before the annex is finalized.

What Is EU GMP Annex 22? Scope, Definitions, and Regulatory Status

Draft Annex 22 is narrowly scoped by design. ECA Academy's gmp-compliance.org describes it as a "supplementary guideline to the EU GMP Guide with specific requirements for Artificial Intelligence" (Source: gmp-compliance.org), applying only to computerised systems in which "Artificial Intelligence models are used in critical applications" with direct impact on patient safety, product quality, or data integrity (Source: health.ec.europa.eu).

The draft's structure is organized as a document map moving through "1. Scope 2. Principles 3. Intended Use 4. Acceptance Criteria" and subsequent operational chapters (Source: health.ec.europa.eu), running to roughly ten chapters plus a glossary.

Definitions matter a great deal in a document this compact, because they determine what would fall inside or outside its proposed scope if the draft is adopted. The glossary defines an **AI system** in language that closely mirrors the EU AI Act's own technology-neutral definition, describing it as "a machine-based system that is designed to operate with" varying levels of autonomy that "may exhibit adaptiveness after deployment" and generates outputs such as predictions, recommendations, or decisions (Source: health.ec.europa.eu). Separately, **machine learning** is defined as "the computational process of optimising the parameters" of a model based on training data (Source: health.ec.europa.eu). These definitions are deliberately broad at the top level (AI system, machine learning) but narrow considerably once the annex reaches its operative scope clauses, which set out proposed expectations for a specific subset of AI technologies discussed later in this report.

Regarding regulatory status, two facts should be kept distinct because they are frequently conflated in secondary commentary. First, Annex 22 is a **draft**; the consultation period during which stakeholders could submit comments has closed, but no adopted final version exists as of August 2026. EudraLex Volume 4 is guidance for interpreting the GMP principles and guidelines established under the applicable EU legislation, rather than a freestanding regulation. Second, the draft was co-authored with PIC/S, whose participating authorities extend beyond the EU's 27 member states; joint drafting demonstrates regulatory cooperation but does not establish how any final text would be adopted or used outside the EU. EMA's continued workshop activity through mid-2026, focused specifically on generative AI, also shows that scope questions remain under consideration, particularly around the treatment of large language models discussed later in this report (Source: ema.europa.eu).

Annex 22 and Annex 11: How the New AI Annex Extends Computerised Systems Requirements

Understanding Annex 22 requires understanding what changed in Annex 11 at the same time, because the two documents were revised as a coordinated package rather than in isolation. Annex 11, the EU's foundational computerised-systems annex, was "initially issued in 2011" and had remained just five pages long for over a decade (Source: bioprocessonline.com). The July 2025 draft revision is dramatically more extensive: trade analysis found "the draft Annex 11 has grown from 5 to 19 pages and the draft Chapter 4 from the original 9 pages to 17 pages," restructured into roughly 17 chapters plus a glossary (Source: gmp-compliance.org). PIC/S summarizes the revised Annex 11 as "mandating that Quality Risk Management principles be comprehensively applied during all steps" of a computerised system's lifecycle (Source: picscheme.org), extending the risk-based validation approach that AI governance under Annex 22 then builds upon.

Annex 22 does not attempt to duplicate this general computerised-systems foundation. Instead, the draft proposes AI-specific expectations that would supplement it, and cross-references Annex 11 directly at several points, stating that it "provides additional guidance to Annex 11 for computerised systems in which AI models are embedded" (Source: health.ec.europa.eu). The clearest example concerns performance acceptance criteria: the draft states that acceptance criteria for an AI model must be set "at least as high as the performance of the process it replaces," explicitly pointing readers to Annex 11 section 2.7 for the underlying computerised-systems requirement being extended (Source: health.ec.europa.eu). In other words, an AI model cannot be validated and deployed simply because it performs adequately in isolation; it must be shown to be at least as reliable as whatever manual process, rules-based system, or earlier model it is replacing.

The practical division of labor between the two documents can be summarized as follows. *Table 1* below compares the two annexes across their core dimensions, and is followed by an interpretive discussion of what the comparison means for compliance teams.

DIMENSION	ANNEX 11 (COMPUTERISED SYSTEMS, REVISED DRAFT)	ANNEX 22 (ARTIFICIAL INTELLIGENCE, NEW DRAFT)
Scope	All GMP-relevant computerised systems, deterministic and AI-based alike (Source: picscheme.org)	Only computerised systems where AI models are used in critical GMP applications
Document length (draft)	Grew from 5 to 19 pages (Source: gmp-compliance.org)	Roughly six pages, deliberately compact (Source: bioprocessonline.com)
Core requirement style	Quality Risk Management applied "during all steps" of the system lifecycle (Source: picscheme.org)	Continuous oversight specific to AI: change control, performance monitoring, human review (Source: picscheme.org)
Technology addressed	General validation, data integrity, audit trails for any computerised system	Explainability, confidence scoring, and scope distinctions for static versus dynamic and deterministic versus probabilistic AI/ML models
Relationship	Foundational; Annex 22 explicitly builds on it	Supplementary guideline to the EU GMP Guide (Source: gmp-compliance.org)

The comparison makes clear that Annex 11 remains the foundation for every computerised system used in GMP manufacturing, including AI-based ones. If adopted, Annex 22 would add a targeted layer of AI-specific controls for systems that cross the draft's "critical" AI-use threshold. A quality unit evaluating a new AI tool should meet applicable Annex 11 and other GMP obligations and may assess the tool against the draft's proposed expectations for explainability, confidence-score logging where applicable, appropriate threshold settings, and model-specific change control. Systems that use AI only in a non-critical, advisory capacity, such as a search or summarization tool that never directly determines a batch disposition decision, may fall outside the draft Annex 22 critical-application scope while remaining subject to applicable GMP requirements.

Core AI-Specific Requirements: Explainability, Confidence Scores, and Human Oversight

The substance of Annex 22 concentrates on a small number of technical controls that have no direct analogue in traditional computerised-systems validation, because they address learned models rather than deterministic, rules-based software. For critical GMP applications, however, the draft itself is limited to learned models with deterministic outputs. A March 2026 peer-reviewed analysis in the PDA Journal of Pharmaceutical Science and Technology characterizes the draft as introducing a structured framework defining expectations for intended use, validation, and lifecycle management (Source: journal.pda.org). Four requirements stand out as the annex's operational core.

Explainability. The draft requires that, during testing of AI models used in critical GMP applications, developers apply interpretability methods such as "feature attribution (e.g. SHAP values or LIME) or visual tools" to understand which input features are driving a model's predictions, and that "a review of these features should be part of the process for approval of test results," meaning explainability output becomes a formal input into the quality sign-off rather than a side artifact (Source: health.ec.europa.eu). SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) are established machine learning techniques that attribute a model's output to specific input variables, letting reviewers judge whether a model is relying on clinically or process-relevant features rather than spurious correlations.

Confidence scoring. The draft states that the system should, **where applicable**, log the model's confidence score for each prediction or classification generated during testing (Source: health.ec.europa.eu). For a very low score, it says that consideration should be given to flagging the outcome as "undecided" rather than making a potentially unreliable prediction; it does not prescribe a universal threshold or require that handling for every model (Source: health.ec.europa.eu).

Human-in-the-loop oversight. For applications involving model types the annex otherwise excludes from critical use, the draft insists that "personnel with adequate qualification and training should always be responsible" for the underlying decision, and that where a model informs a human decision rather than acting autonomously, "this may imply a consistent review and/or test of every output from the model" depending on the criticality of the application, rather than a sampling-based approach (Source: health.ec.europa.eu).

Data and testing independence. To prevent evaluation bias, the draft requires effective independence controls between the personnel and data used to train a model and those used to test it, including access controls, audit trails, and what the draft calls a "colleague who has not had this access (4-eyes principle)" for reviewing test outcomes (Source: health.ec.europa.eu). The four-eyes principle, long familiar in GMP quality control for physical checks such as weighing and dispensing, is here extended specifically to prevent a single individual from both curating a model's training data and independently certifying its test results.

Underpinning all four requirements is a stated risk-based philosophy: activities under Annex 22 should be "implemented based on the risk to patient safety, product quality and data integrity" (Source: health.ec.europa.eu). In practice, the depth of explainability documentation and the frequency of human review should be designed according to how directly a given AI application touches patient safety, product quality, and data integrity. The draft separately calls for confidence-score logging where applicable and appropriate threshold settings, with an undecided outcome to be considered when a score is very low.

Static vs Dynamic AI Models: Scope Boundaries and the Generative AI Exclusion

Perhaps the single most consequential design decision in draft Annex 22 is where it proposes a scope boundary among different classes of AI model, because that boundary would determine which technologies are covered by the proposed guidance versus which the draft says should not be used in critical GMP applications.

The annex's glossary defines a **static model**, which it also calls a "frozen model," as "a model where all parameters have been finally set" and that does "not adapt their performance during use by incorporating new data" once validated and deployed (Source: health.ec.europa.eu). By contrast, PIC/S's own summary of the draft describes **dynamic models** as those "which continuously and automatically learn and adapt performance during use" (Source: picscheme.org). The draft is unambiguous about its critical-application boundary: it governs static models with deterministic outputs. Dynamic-model behavior "is not covered by this document, and should not be used in critical GMP" applications at all; the draft likewise excludes models with probabilistic outputs, which may produce different outputs from identical inputs, from critical GMP applications (Source: health.ec.europa.eu).

The draft applies an equally strict exclusion to a second category of technology. PIC/S's own consultation documentation states that "the document does not apply to Generative AI and Large Language Models (LLM)" (Source: picscheme.org), a scope limitation the underlying European Commission draft text restates in near-identical terms. In practice this means the current draft is written for a relatively conservative and well-established class of technology, predictive and classification models such as computer-vision quality inspection systems, predictive-maintenance models, or spectroscopic anomaly detectors, all of which can plausibly be frozen and validated as static artifacts. It is not, at least not yet, a governance framework for the generative and conversational AI tools that have become widespread in other business functions.

Table 2 below summarizes the distinction, followed by a discussion of its practical and forward-looking implications.

MODEL CATEGORY	DEFINITION PER DRAFT ANNEX 22	REGULATORY TREATMENT UNDER THE CURRENT DRAFT
Static (frozen) model	"A model where all parameters have been finally set" (Source: health.ec.europa.eu)	Covered; subject to explainability and change-control provisions, with confidence-score logging where applicable and undecided-outcome handling to be considered for very low scores
Dynamic (adaptive) model	Continuously and automatically learns and adapts performance during use (Source: picscheme.org)	Not covered; should not be used in critical GMP applications (Source: health.ec.europa.eu)
Generative AI / LLMs	Not separately defined; addressed as an exclusion category (Source: picscheme.org)	Out of scope and should not be used in critical GMP applications. For non-critical GMP applications, qualified and trained personnel should be responsible for ensuring outputs are suitable for the intended use; human-in-the-loop principles may be considered where applicable.

This exclusion has not gone unchallenged. EMA's own account of the 2025 consultation responses notes that feedback "suggested support for potentially enabling the use of technologies such as generative AI" in medicines manufacturing, despite the draft's current position (Source: ema.europa.eu). That feedback appears to have driven EMA's June to July 2026 expert workshop, whose stated purpose was to help shape a risk-based approach to generative AI specifically, with workshop questions asking experts to apply quality risk management to "adaptive" and "probabilistic" AI model lifecycles, "taking into account ICH Q9R1 principles" on quality risk management (Source: ema.europa.eu). The workshop's identified guardrail principles for any future expansion of scope center on "data governance, model evaluation, transparency, accountability, and human oversight" (Source: ema.europa.eu), the same core themes as the static-model requirements already in the draft, extended conceptually to a harder problem: models whose behavior can itself evolve after validation. Manufacturers should monitor the current exclusions for generative AI, dynamic models, and probabilistic-output models as the drafting process continues. EMA's workshop shows that possible controls for generative AI remain under consideration; it does not establish the scope or requirements of any final text.

AI Model Lifecycle Management: Validation, Change Control, and Drift Monitoring

Beyond the point-in-time testing requirements covered above, Annex 22 treats AI model governance as a continuing lifecycle obligation rather than a one-time validation exercise, an emphasis that PIC/S's own summary captures directly: "Annex 22 foresees a continuous oversight of AI systems, including change control" and model performance monitoring (Source: picscheme.org).

Change control before deployment. Once a model, its surrounding system, and the manufacturing process it automates have been tested and accepted, the draft requires that they "be put under change control before it is deployed in operation," meaning any subsequent modification (retraining, parameter adjustment, infrastructure change) must trigger a formal evaluation of whether retesting is required before the change goes live (Source: health.ec.europa.eu). This treats a validated static model the same way traditional GMP treats a validated analytical method or a qualified piece of equipment: changes are not forbidden, but they cannot happen silently.

Drift monitoring. Because even a frozen model's real-world performance can degrade if the data it encounters in production diverges from the data it was trained and tested on, the draft requires ongoing surveillance of that divergence: "metrics should be defined for monitoring any drift in the input data" relative to the model's original sample space (Source: health.ec.europa.eu). This is a distinctly AI-specific concept with no direct equivalent for conventional software validation; a rules-based system does not silently become less accurate as raw material lots or sensor calibrations shift over time in the way a trained statistical model can.

Risk-based scaling. As previously noted, the draft proposes that activities across the model lifecycle, testing depth, review frequency, and monitoring intensity be implemented according to risk to patient safety, product quality, and data integrity. PIC/S's parallel Annex 15 (Qualification and Validation) revision work considered revision of ICH Q9(R1) on quality risk management; the joint concept-paper consultation closed on 9 April 2026, indicating that risk-management methodology across the GMP guideline family is under review (Source: picscheme.org).

Taken together, these lifecycle provisions mean that an organization cannot treat AI model validation as complete once a model passes initial acceptance testing. Quality systems need mechanisms to detect performance degradation, mechanisms to control and re-evaluate model changes, and governance that connects both back to a documented, risk-tiered rationale, effectively converting model management from a project into an ongoing operational discipline analogous to preventive maintenance or continued process verification, a point the PDA Journal analysis frames explicitly in terms of continued process verification (Source: journal.pda.org).

The Global Regulatory Context: PIC/S, FDA, ICH, WHO, MHRA, and TGA

Draft Annex 22 did not emerge in isolation. It is part of a broader, uneven pattern of global regulatory engagement with AI in pharmaceutical manufacturing, in which some bodies have moved to detailed proposed text, others have issued discussion papers or reflection papers, others are actively revising adjacent risk-management guidelines, and at least one has concluded that formal guidance would currently be premature.

PIC/S, the international scheme whose participating authorities work toward harmonized GMP inspection standards, co-drafted Annex 22 directly with EMA and ran the identical three-month stakeholder consultation in parallel with the EU process (Source: picscheme.org). Joint drafting does not itself establish adoption, applicability, or enforcement outside the EU; any effect in a participating authority would depend on the relevant PIC/S and national processes. PIC/S's Annex 15 (Qualification and Validation) revision work also considered possible revision of ICH Q9(R1) on quality risk management; the related joint concept-paper consultation closed on 9 April 2026 (Source: picscheme.org).

The FDA has taken a more incremental, discussion-paper-led approach rather than drafting binding regulatory text. The agency's Center for Drug Evaluation and Research (CDER) published a discussion paper titled "Artificial Intelligence in Drug Manufacturing" in February 2023 (Source: fda.gov), and CDER's Framework for Regulatory Advanced Manufacturing Evaluation (FRAME) initiative subsequently "prioritized artificial intelligence (AI) as a technology that has the potential to advance pharmaceutical manufacturing" (Source: link.springer.com). A subsequent FDA/PQRI public workshop on 26 and 27 September 2023 found that "approximately one-third of attendees are actively utilizing AI for data analysis, process development" and related manufacturing work, even before any binding guidance existed (Source: link.springer.com). FDA's Emerging Technology Program, run out of CDER's Office of Pharmaceutical Quality since 2014, has separately supported approvals using innovative manufacturing methods, including advanced modeling approaches relevant to AI-enabled processes (Source: fda.gov).

The FDA's enforcement posture crystallized in April 2026 with what industry commentary describes as its first cGMP warning letter to explicitly name inappropriate AI use as a violation, discussed at greater length in the Case Studies section below. The agency cited a manufacturer for having "used AI to create drug product specifications, procedures, and master production" records without adequate quality oversight, noting that "overreliance on artificial intelligence for your drug manufacturing operations was also documented during the inspection" (Source: fda.gov) (Source: fda.gov). Its core enforcement principle mirrors the EU draft's human-in-the-loop philosophy: "any output or recommendations from an AI agent must be reviewed and cleared by an" authorized quality representative before use in GMP documentation (Source: fda.gov).

ICH (the International Council for Harmonisation) has moved to formally place AI-based manufacturing technology on its future guideline agenda. Its Assembly endorsed a Reflection Paper on 8 October 2025 that lists "process modelling, including artificial intelligence (AI)-based models" as an example of the advanced manufacturing technologies future ICH guideline work should address (Source: admin.ich.org), signaling that AI-specific manufacturing guidance may eventually be harmonized globally through ICH rather than remaining a purely EU or PIC/S initiative.

WHO, by contrast, has been notably more cautious. At a June 2023 GMP consultation, the organization concluded "it seems too early to draft a guideline on the topic as there is not enough" practical implementation information yet available from manufacturers and regulators (Source: cdn.who.int). WHO's foundational AI-health work instead runs through "a Focus Group on AI for Health (FG-AI4H), jointly with ITU," created in 2018 to develop benchmarking and standardization methodology (Source: cdn.who.int), while industry representatives at the same meeting noted that even the International Coalition of Medicines Regulatory Authorities' 2021 horizon-scanning report on AI, described as having "issued the Horizons scanning assessment report," had not addressed manufacturing-specific applications at all (Source: cdn.who.int).

The UK's MHRA and Australia's TGA have, so far, addressed AI regulation primarily through the lens of AI as a medical device (AIaMD) rather than AI within GMP manufacturing specifically. MHRA's flagship initiative, the AI Airlock regulatory sandbox, ran its pilot phase "between April 2024 and March 2025" and is explicitly scoped to AI-enabled medical device software rather than manufacturing computerised systems, according to both the summary report page (Source: gov.uk) and the underlying full programme report document (Source: assets.publishing.service.gov.uk). Similarly, the TGA regulates AI-enabled software primarily by its intended purpose, noting that "developers of AI-enabled medical device software may meet the definition of a manufacturer or sponsor" under the Therapeutic Goods Act, a device-centric framework that does not yet extend a parallel structure to AI used in drug manufacturing GMP processes (Source: tga.gov.au).

Separately from formal guideline text, **EMA** itself published a broader "Reflection paper on the use of artificial intelligence in the medicinal product lifecycle," adopted 9 September 2024, which calls for organizations to "integrate data science competence with the respective fields within medicines development, manufacturing and pharmacovigilance" (Source: ema.europa.eu) and states that "SOPs implementing GxP principles on data and algorithm governance should be extended" to cover AI and machine learning systems used in high-impact settings (Source: ema.europa.eu). This reflection paper predates and conceptually frames Annex 22, extending the same governance philosophy, data governance, algorithm oversight, GxP-aligned SOPs, across the entire medicinal product lifecycle rather than manufacturing alone.

Table 3 below summarizes each body's current posture, followed by a synthesis of what the pattern means for globally operating manufacturers.

BODY	PRIMARY INSTRUMENT	STATUS AS OF AUGUST 2026
EU (EMA/European Commission)	Draft Annex 22, revised Annex 11, revised Chapter 4	Consultation closed; further expert workshop held mid-2026; not yet finalized (Source: health.ec.europa.eu)
PIC/S	Joint Annex 22 co-authorship; parallel Annex 15/ICH Q9(R1) review	Consultation closed; ICH Q9(R1) revision window open to 9 April 2026 (Source: picscheme.org)
FDA (US)	Discussion papers, FRAME initiative, enforcement via warning letters	No binding AI-manufacturing rule; active enforcement precedent set April 2026 (Source: raps.org)
ICH	Reflection Paper on advanced manufacturing technologies	Endorsed October 2025; names AI-based process modelling as future guideline scope (Source: admin.ich.org)
WHO	None specific to GMP manufacturing	Concluded in 2023 that formal guidance was premature (Source: cdn.who.int)
MHRA (UK) / TGA (Australia)	AI Airlock sandbox / device-centric AI rules	Scoped to AI as a medical device, not GMP manufacturing (Source: gov.uk) (Source: tga.gov.au)

The overall pattern is one of the EU and PIC/S leading with concrete draft regulatory text, ICH signaling intent to harmonize globally, and the FDA regulating through enforcement precedent and discussion papers rather than binding rules. WHO has described manufacturing-specific guidance as premature, while MHRA's cited AI programme is device-focused and TGA is participating in the PIC/S Annex 22 revision process relevant to Australian GMP manufacturers. For multinational manufacturers, the practical implication is to monitor their applicable authorities and maintain controls that satisfy existing GMP obligations; the available sources do not establish that a final Annex 22 would become a global benchmark or be referenced informally by other jurisdictions' inspectors.

Implementation Checklist: Preparing for Potential Annex 22 Expectations

Because Annex 22 remains in draft form, it does not add adopted Annex 22 guidance; manufacturers nevertheless remain subject to the applicable GMP legal requirements and related quality-system obligations. The FDA's April 2026 enforcement action against AI overreliance in manufacturing documentation illustrates that existing GMP controls remain relevant to AI-enabled processes (Source: raps.org). Quality and manufacturing leaders may use the following checklist to assess readiness against the July 2025 draft's proposed expectations; it does not substitute for the final text or applicable GMP obligations.

- **Inventory existing AI use.** Catalog every computerised system in the facility that incorporates a machine learning or AI component, and classify each by whether it touches a "critical" GMP application affecting patient safety, product quality, or data integrity (Source: health.ec.europa.eu).
- **Classify each model by adaptation and output behavior.** For critical use under the current draft, a model must be static (frozen) and provide deterministic outputs; dynamic and probabilistic-output models should not be used in critical GMP applications.
- **Flag generative AI and LLM use separately.** Any generative AI or LLM tool touching GMP-critical decisions should be treated as out of scope under the current draft; if used in non-critical GMP applications, qualified and trained personnel should ensure outputs are suitable for the intended use (Source: picscheme.org).
- **Build explainability documentation into validation protocols**, incorporating feature-attribution analysis such as SHAP or LIME into model test plans and routing that analysis through formal quality review before test-result approval.

- **Assess confidence-score logging and undecided-outcome handling.** Where applicable, log confidence scores for predictions or classifications; for very low scores, consider whether an undecided outcome and an appropriate documented follow-up process are warranted.
- **Establish AI-specific change control**, bringing each validated model under formal change control before go-live, with defined triggers for re-validation when training data, parameters, or infrastructure change.
- **Deploy drift monitoring**, defining and continuously tracking input-data drift metrics against the model's original training and test data envelope.
- **Separate training and test personnel and data**, applying access controls and a four-eyes principle so the individuals who curate training data are not the same individuals who independently certify test outcomes.
- **Benchmark against the replaced process**, setting acceptance criteria at least as strong as the performance of whatever manual or legacy process the AI model is replacing.
- **Document human-in-the-loop accountability where the draft calls for it.** Where a model informs a human operator's decision and testing effort has been diminished, document the operator's responsibility and keep records of that process. The extent of consistent review or testing of outputs depends on process criticality and the model's level of testing. For generative AI or LLMs in non-critical GMP applications, ensure qualified and trained personnel are responsible for output suitability.
- **Align risk tiers to Quality Risk Management**, scaling the intensity of validation, monitoring, and review to the model's potential impact, consistent with the QRM principles both Annex 22 and the revised Annex 11 emphasize (Source: picscheme.org).
- **Track the finalization timeline**, monitoring official regulatory consultation pages and PIC/S announcements for the adopted final text, since specific thresholds and wording may shift from the July 2025 draft.

Life-sciences advisory practices that specialize in regulatory compliance and AI adoption typically frame this kind of checklist work as a distinct advisory service line, offering "strategic guidance on digital transformation, AI adoption, and technology roadmapping" for organizations navigating exactly this kind of pre-finalization uncertainty (Source: intuitionlabs.ai). Whether an organization runs this checklist internally through its quality and validation functions or engages outside regulatory and AI advisory expertise, the underlying task is the same: build the explainability, confidence-scoring, change-control, and human-oversight infrastructure now, so that whichever final wording Annex 22 adopts, the organization is not starting from zero.

Data Analysis and Evidence

Quantifying the AI-in-pharmaceutical-manufacturing landscape requires triangulating across several sources that use different methodologies, scopes, and forecast horizons, and the resulting figures do not always agree, a discrepancy worth stating openly rather than collapsing into a single misleadingly precise number.

Market size estimates. Fortune Business Insights projects the global AI-in-pharma-manufacturing market specifically to grow from $1.64 billion in 2026 to $12.90 billion by 2034 (Source: fortunebusinessinsights.com). MarketsandMarkets, covering the broader "AI in life science" category rather than manufacturing specifically, projects growth "from USD 21.58 billion in 2026 to USD 69.34 billion by 2031, at a CAGR of 26.3%" (Source: marketsandmarkets.com). BCC Research, covering the "AI in pharmaceuticals" category, estimates growth "from $3.8 billion in 2025 to reach $15.2 billion by the end of 2030" (Source: bccresearch.com). The wide spread among these three figures, ranging from roughly $1.6 billion to $21.6 billion for overlapping years, mainly reflects differing category definitions (manufacturing alone versus the entire life-science AI market) rather than genuine disagreement about growth trajectory; all three independently forecast compound annual growth in the 25 to 32% range, a consistency that itself corroborates the broader adoption trend documented by Rockwell Automation's survey of life-sciences manufacturers (Source: rockwellautomation.com).

Industry adoption surveys. The ISPE Pharma 4.0 Survey, a longitudinal study the International Society for Pharmaceutical Engineering has run since 2017, draws on cumulative "data gathered from nearly 3,200 respondents across the life sciences sector" through its most recent editions (Source: ispe.org); its 7th edition analyzed "the 2023 survey sample, consisting of 418 respondents" spanning a geographically diverse base (Source: ispe.org). Separately, the 2025 Rockwell Automation survey covering 143 life-sciences manufacturing leaders across 15 countries found that "95% of life sciences manufacturers are using or evaluating smart technology" including AI (Source: rockwellautomation.com), with the leading reported use cases being that "companies use AI to improve quality (53%), streamline operations (50%) and strengthen cybersecurity (48%)" (Source: rockwellautomation.com).

Realized versus projected benefits. A 2025 Deloitte Center for Health Solutions survey of 103 biopharma executives found quality-control lab modernization already delivering measurable, realized gains: "50% of survey respondents reported fewer errors and deviations, 45% noted improved compliance," and 43% observed shorter testing timelines (Source: deloitte.com). The same executives offered a distinctly more speculative set of forward-looking figures that should be read as forecasts, not realized results: they project further digitization "could lead to a 20% to 50% reduction in compliance issues, a 15% to 30% decrease in operational costs" (Source: deloitte.com). Separate Deloitte modeling estimates that a top-ten

biopharmaceutical company could capture "between $5-7 Bn of peak value by scaling the use of AI over 5 years," with manufacturing and supply chain applications representing an estimated 15 to 25% of that total value pool (Source: deloitte.com), a figure broadly consistent with the operational gains that named manufacturers such as Merck and Roche report from their own deployments, discussed in the Case Studies section below.

Enforcement and data-integrity baseline. Because Annex 22's requirements are ultimately about controlling risk to data integrity, it is useful to benchmark the current enforcement environment those controls sit within. A peer-reviewed study conducting "full-enumeration analysis of 1766 FDA Warning Letters issued between 2016 and 2023" and reclassifying violations under the ALCOA/ALCOA+ data-integrity framework found that "the average number of DI violations per company increased in 2023" relative to prior years (Source: pubmed.ncbi.nlm.nih.gov) (Source: pubmed.ncbi.nlm.nih.gov). More recent enforcement data show the same upward pressure: "FDA issued a total of 190 warning letters to drug and biologics manufacturers in Fiscal Year 2024" (Source: pharmaceuticalonline.com), and within that total, "FDA issued 105 warning letters to human drug manufacturing sites, marking the highest number" of such letters in five years (Source: raps.org). Separate reporting on FDA's FY2024 drug quality report found that "in FY2024, more than 62% of drug quality assurance inspections were at foreign sites," an all-time high for the agency (Source: raps.org).

Read together, this data indicates that use or evaluation of smart-manufacturing technology is widespread among the surveyed life-sciences manufacturers, while the survey separately reports AI use cases. It also shows measurable quality gains, though not necessarily order-of-magnitude transformation; market forecasts that agree directionally on strong double-digit growth despite differing category definitions; and a regulatory enforcement backdrop that gives urgency to explainability, conditional confidence provisions, and change-control disciplines proposed in Annex 22.

Case Studies and Real-World Examples

The following six examples illustrate reported AI use in pharmaceutical manufacturing and regulatory response when use goes wrong. They include an FDA warning letter and company- or vendor-published deployment accounts; the performance claims in those accounts should be read as reported outcomes, not independent verification.

FDA Warning Letter: Purolea Cosmetics Lab (April 2026)

In April 2026, the FDA issued a warning letter to Purolea Cosmetics Lab that industry legal commentary describes as "the first time this has appeared as a named violation in an FDA warning letter," referring specifically to inappropriate reliance on AI in drug manufacturing documentation (Source: kneat.com). FDA investigators found the firm had "used AI to create drug product specifications, procedures, and master production" records without the quality unit review CGMP requires (Source: fda.gov). RAPS's independent coverage confirms the letter cited "excessive reliance on artificial intelligence (AI) to create drug specifications, procedures, and production records" as the underlying violation (Source: raps.org). A particularly telling detail from the warning letter: when investigators asked why the firm had not completed a required process validation step, the company's response was that its AI agent had "never told you it was required" (Source: fda.gov), an answer FDA plainly rejected as inadequate, since the letter's corrective expectation is unambiguous: "any output or recommendations from an AI agent must be reviewed and cleared by an" authorized, qualified human before it can support a regulatory decision (Source: fda.gov). This case functions as a real-world preview of the human-in-the-loop principle draft Annex 22 seeks to formalize on the EU side, arrived at independently through US enforcement.

Novo Nordisk: Computer Vision for Insulin Manufacturing

An AWS-published case study reports that Novo Nordisk, a manufacturer that "supplies nearly 50 percent of the world's insulin," deployed AWS-based computer vision paired with machine learning to automate quality tasks including cartridge counting and agar-plate anomaly detection on its manufacturing lines (Source: aws.amazon.com). The account describes the deployment as "computer vision combined with machine learning (ML) to automate key tasks" using a robotic arm, camera rig, edge-device inference, and an Amazon SageMaker pipeline covering training, deployment, and ongoing monitoring (Source: aws.amazon.com). This vendor-published account is not independent verification. It describes the type of well-bounded computer-vision quality-inspection application that could fall within draft Annex 22's critical-use scope only if the deployed model is static and provides deterministic outputs.

Merck (MSD): HawkAVI Automatic Visual Inspection Platform

Merck & Co. (known as MSD outside the United States and Canada) built an AI/ML platform called "HawkAVI" on AWS to analyze automatic visual inspection reject images, aiming to distinguish genuine product defects from false rejects. In an AWS post guest-authored by MSD employees, the company reported that the platform reduced false-reject rates by 50% across various product lines (Source: aws.amazon.com). This is a company-

reported outcome, not an independently verified benchmark. The same post reports that the Manufacturing Leadership Council awarded the platform a 2021 Manufacturing Leadership Award (Source: aws.amazon.com).

Novartis: AI Insight Centers and CAR-T Manufacturing

An AWS blog post reports that Novartis, "operating more than 60 manufacturing sites" globally, partnered with AWS beginning in 2019 to build so-called "Insight Centers" applying machine learning to predict site performance issues and enable computer-vision-based line clearance checks (Source: aws.amazon.com), noting at the time that "Novartis is currently using Amazon SageMaker to build a computer vision-based model" for that purpose (Source: aws.amazon.com). This vendor-published account is not independent verification. Separately, Novartis signed a five-year AI research alliance with Microsoft in October 2019 that specifically included applying AI to improve "the manufacturing of CAR-T cells, a time-consuming and resource intensive" process for personalized cell therapy production (Source: healthcaredive.com), illustrating AI's extension beyond conventional small-molecule and biologics manufacturing into advanced cell and gene therapy production.

Roche/Genentech: Predictive Yield Modeling and Deviation Management

According to reporting on Roche's technical operations, the company "has already scaled AI deployments, delivering clear time and cost savings" across its manufacturing network, moving past isolated pilots (Source: bioprocessonline.com). Roche's AI-powered predictive applications for biologics titer and yield forecasting, along with critical quality attribute prediction, reportedly delivered a "5% to 10% boost in yields and enhanced target product quality" (Source: bioprocessonline.com). Roche additionally deployed AI-driven scheduling algorithms across its global manufacturing network to optimize product sequencing during planned downtimes, reporting "productivity gains of 5% to 10% on constrained assets" (Source: bioprocessonline.com). On the quality side, "Roche Pharma's approach to manufacturing deviation management through the innovative use" of generative AI with retrieval-augmented generation helps teams identify and analyze historically similar deviations faster (Source: bioprocessonline.com). If such a tool were used in a critical GMP application, the current draft says generative AI and LLMs should not be used. For a non-critical GMP application, personnel with adequate qualification and training should be responsible for ensuring output suitability, while human-in-the-loop principles may be considered where applicable (Source: health.ec.europa.eu).

GSK: Digital Control Rooms at Upper Merion, Pennsylvania

Following a $120 million investment in its Upper Merion, Pennsylvania site, GSK deployed AI and machine learning mathematical models where "digital mathematical models are enabling teams to predict and improve the" reliability of manufacturing equipment, reducing rejected product batches (Source: gsk.com). The site's digital control room, combined with AI-enabled exception-based batch record review, was able to "reduce the number of pages to be reviewed from 100s per batch" down to a small fraction requiring manual quality attention, roughly 20 to 30 pages (Source: gsk.com). This exception-based review model illustrates a possible approach to confidence-based escalation. Draft Annex 22 conditionally addresses confidence-score logging and says an undecided outcome should be considered for very low scores; it does not prescribe this specific operating model.

Taken together, these six cases span the full spectrum the rest of this report has discussed: an enforcement action illustrating the cost of skipping human oversight, several computer-vision deployments that could be within draft Annex 22's intended scope if their models are static and deterministic, and a generative-AI deviation-management use case that is excluded from critical GMP applications under the current draft. No single case demonstrates full formal compliance with Annex 22 because it remains an unadopted draft; collectively, however, the cases show that operational patterns addressed by the draft—explainability, confidence-based escalation, human accountability, and change-controlled deployment—are already emerging in industry practice.

Implications and Future Directions

Several forward-looking dynamics are likely to shape how Annex 22 evolves and how manufacturers should plan around it.

First, the generative AI and LLM exclusion is likely to narrow rather than remain permanent. EMA's own account of the 2025 consultation notes stakeholder feedback "suggested support for potentially enabling the use of technologies such as generative AI" (Source: ema.europa.eu), and the agency's mid-2026 expert workshop was convened specifically to work out a risk-based path toward that expansion, referencing ICH Q9(R1) quality risk management principles as the analytical foundation (Source: ema.europa.eu). Organizations already piloting generative AI in adjacent, non-critical manufacturing functions, such as Roche's own deviation-management application described above (Source: bioprocessonline.com), should treat the current draft's human-in-the-loop requirement as a durable design principle even as the formal exclusion itself may eventually loosen.

Second, harmonization pressure is building from multiple directions simultaneously. PIC/S's joint authorship of Annex 22 demonstrates cooperation with EU authorities, but it does not establish that any final text will be adopted, applied, or used as an informal benchmark outside the EU. ICH's October 2025 Reflection Paper signals that AI-based process modelling may eventually receive its own globally harmonized ICH guideline rather than remaining a purely EU initiative (Source: admin.ich.org). PIC/S's parallel work reconsidering ICH Q9(R1) on quality risk management included a joint concept-paper consultation that closed on 9 April 2026; any resulting revision could affect the risk-management vocabulary relevant to Annex 22 (Source: picscheme.org).

Third, enforcement is arriving faster than formal AI-specific guidance in at least one major jurisdiction. The FDA's April 2026 warning letter to Purolea Cosmetics Lab demonstrates that regulators will act against inappropriate AI reliance in manufacturing under existing CGMP authority, without waiting for final AI-specific GMP guidance (Source: fda.gov). Manufacturers should not treat the absence of finalized AI-specific guidance as license to defer good-practice controls; the underlying quality obligations that AI misuse can violate, adequate specifications, validated processes, and documented human review, already exist and are already being enforced.

Fourth, the gap between adoption and governance maturity remains wide. With roughly 95% of surveyed life-sciences manufacturers already using or evaluating AI-adjacent smart-manufacturing technology (Source: rockwellautomation.com), but no jurisdiction yet operating under finalized, binding, AI-specific GMP text, most organizations are operating in a compliance gray zone: technically unregulated by AI-specific rules, but fully subject to existing GMP, data-integrity, and quality-system obligations that AI misuse can still violate, as the FDA warning letter illustrates. Life-sciences consultancies with combined regulatory and AI implementation expertise describe this gap as their core value proposition, noting that they "understand the unique regulatory landscape, data complexities, and business drivers of the life sciences sector" precisely because organizations need help translating draft regulatory principles into working quality-system controls before those principles become binding law (Source: intuitionlabs.ai). Whichever route an organization takes, internal quality build-out or external advisory support, the direction of travel is clear: explainability, confidence-based escalation, drift monitoring, and human accountability are converging as the baseline expectations for AI in regulated manufacturing, regardless of exactly when or in what final form Annex 22 is formally adopted.

Frequently Asked Questions (FAQs)

What is EU GMP Annex 22? Annex 22 is a proposed new annex to EudraLex Volume 4, the EU's compiled Good Manufacturing Practice guidance, setting out proposed expectations for the use of AI in GMP-regulated pharmaceutical manufacturing. It is a "supplementary guideline to the EU GMP Guide with specific requirements for Artificial Intelligence" (Source: gmp-compliance.org) and was jointly drafted by EMA and PIC/S (Source: picscheme.org).

What are the main AI-specific requirements in Annex 22? The draft addresses explainability techniques such as SHAP or LIME during model testing (Source: health.ec.europa.eu), confidence-score logging where applicable, consideration of an "undecided" outcome when a score is very low, human accountability for excluded model types used in non-critical GMP applications, pre-deployment change control, and drift monitoring. PIC/S summarizes these expectations as "continuous oversight of AI systems, including change control" (Source: picscheme.org).

How does Annex 22 differ from Annex 11? Annex 11 covers all GMP-relevant computerised systems generally, applying Quality Risk Management "during all steps" of a system's lifecycle (Source: picscheme.org), while Annex 22 adds AI-specific requirements, explainability, confidence scoring, and the static/dynamic model distinction, layered specifically on top of that foundation for systems where AI models are used in critical applications.

When does EU GMP Annex 22 come into force? As of this report's 5 August 2026 publication date, Annex 22 remains a draft. The joint EU/PIC/S public consultation ran from 7 July 2025 to 7 October 2025 and has since closed, with the official European Commission page confirming "the response period for this consultation has ended" (Source: health.ec.europa.eu). No finalized or adopted Annex 22 guidance has yet been published. Existing GMP obligations continue to arise from applicable EU legislation; EudraLex Volume 4 provides guidance for interpreting GMP principles and guidelines. EMA convened a further expert workshop on generative AI on 30 June and 1 July 2026 (Source: ema.europa.eu), indicating finalization is still pending.

Does Annex 22 apply to generative AI or large language models? No, not for critical GMP applications under the current draft. PIC/S's summary states "the document does not apply to Generative AI and Large Language Models (LLM)" (Source: picscheme.org). EMA's mid-2026 workshop activity indicates that possible controls for generative AI remain under consideration, but does not establish the scope of a future revision.

What is the difference between static and dynamic AI models under Annex 22? A static (or "frozen") model has "all parameters ... finally set" and does not learn further after deployment; a dynamic model "continuously and automatically" learns and adapts during use (Source: picscheme.org). Only static models with deterministic outputs are covered by the current draft for critical applications; dynamic models and probabilistic-output models should not be used in critical GMP contexts.

What is PIC/S's role and guidance on AI in pharmaceutical manufacturing? PIC/S co-drafted Annex 22 alongside EMA's Inspectors Working Group and participated in the joint consultation process. This cooperation does not itself establish adoption, applicability, or enforcement outside the EU; any effect in a PIC/S participating authority depends on the relevant PIC/S and national processes. PIC/S also consulted on a related Annex 15 concept paper that considered revision of ICH Q9(R1) on quality risk management; that consultation closed on 9 April 2026 (Source: picscheme.org).

How should manufacturers approach AI model validation under Annex 22 today? Even before finalization, manufacturers should build explainability documentation, confidence-score logging, change control, and drift monitoring into AI validation protocols now, using the checklist in this report as a starting framework, since the FDA has already shown it will enforce against inadequate AI oversight under existing CGMP authority (Source: raps.org).

Conclusion

Draft EU GMP Annex 22 is among the most detailed proposed AI-specific GMP texts for pharmaceutical manufacturing, setting out proposed expectations for explainability, conditional confidence provisions, human oversight, change control, and drift monitoring within a compact, risk-based framework that would supplement the existing Annex 11 computerised-systems regime if adopted (Source: picscheme.org). Its proposed scope is deliberately conservative: the draft covers static, frozen models with deterministic outputs used in critical GMP applications, while stating that dynamic, continuously adapting models and probabilistic-output models should not be used in critical GMP applications and excluding generative AI and large language models from the document's scope (Source: picscheme.org), pending further work that EMA's mid-2026 expert workshop suggests is already underway (Source: ema.europa.eu). The annex's joint authorship with PIC/S demonstrates regulatory cooperation; it does not itself establish adoption, applicability, or enforcement beyond the EU (Source: picscheme.org), and its core principles, explainability, confidence-based escalation, human accountability, and continuous monitoring, are echoed independently in FDA enforcement action (Source: fda.gov), EMA's broader medicinal-product lifecycle reflection paper (Source: ema.europa.eu), and ICH's advanced-manufacturing reflection paper (Source: admin.ich.org), suggesting a genuine convergence of regulatory thinking rather than an isolated European initiative.

For manufacturing, quality, and regulatory affairs leaders, the practical takeaway is not to wait for a final published text before acting. In Rockwell Automation's survey, 95% of life-sciences manufacturers reported using or evaluating smart technology, while the same survey separately identified AI use cases (Source: rockwellautomation.com); AWS-hosted company accounts describe deployments at manufacturers such as Novo Nordisk and Merck, with MSD reporting a reduction in false-reject rates (Source: aws.amazon.com). These accounts are not independently verified evidence of performance. At least one regulator has already shown it will treat inadequate AI oversight as a CGMP violation under existing authority (Source: raps.org). Organizations preparing for possible final Annex 22 guidance can build explainability, confidence-score logging where applicable, appropriate threshold settings, change control, and human-in-the-loop accountability into their AI-enabled manufacturing systems, while continuing to assess those controls against applicable GMP obligations and any final adopted text.

Tags: eu gmp annex 22, annex 22 artificial intelligence, gmp compliance, annex 11 computerised systems, pic/s guidance, ai validation gmp manufacturing, static vs dynamic ai models, pharmaceutical manufacturing ai

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

Leading Claude & ChatGPT AI Enablement and Training: Help biotech and pharmaceutical teams adopt Claude and ChatGPT through practical training, governed workflows, use-case development, and implementation designed for regulated environments.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at intuitionlabs.ai/contact for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.