AI agent governance in pharma R&D: which standards and guidance apply?

Eight documents and initiatives mapped by scope, status and whether they reach discovery research or GxP work. None is binding law. Statuses checked against each issuing body's page on 30 October 2026 (ISO/IEC 42001: ISO page as archived 6 October 2026, because iso.org blocks automated checks). Drafts and initiatives can change: the linked pages are the authority.

Document (issuing body)	What it covers	Status	Discovery R&D	GxP work	Primary source
Pistoia Alliance agentic AI initiative (Pistoia Alliance)	A pre-competitive project to develop a life sciences and healthcare agent-to-agent communication protocol and an AI agent standard. Planned outputs named in the announcement: whitepapers, guidelines, reference implementations and scientific publications.	Initiative. Announced 4 September 2025 with seed funding from Genentech and a call for more sponsors. The announcement describes planned outputs; no published standard found on the Alliance's site (checked 30 October 2026).	Yes, by design: it sits under the Alliance's priority to harness AI to expedite R&D.	No regulatory standing. Useful as an emerging industry reference, not as a GxP requirement.	https://pistoiaalliance.org/ai/pistoia-alliance-unveils-agentic-ai-initiative-and-seeks-industry-funding-to-drive-safe-adoption/
Considerations for the Use of AI to Support Regulatory Decision-Making for Drug and Biological Products (US FDA)	AI models that produce information or data to support regulatory decisions on drug safety, effectiveness or quality. Proposes a risk-based credibility assessment framework tied to the model's context of use.	Draft, January 2025, docket FDA-2024-D-4689; marked draft and not for implementation (checked 30 October 2026).	No. The draft says it does not address AI used in drug discovery, or for operational efficiencies that do not affect patient safety, drug quality or study reliability.	Yes, where AI output supports a regulatory decision in the nonclinical, clinical, postmarketing or manufacturing phases.	https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
Guiding principles of good AI practice in drug development (European Medicines Agency and US FDA, jointly)	Ten high-level principles for AI used to generate or analyse evidence across the drug product life cycle, including nonclinical, clinical, post-marketing and manufacturing phases: for example a risk-based approach, a clear context of use, adherence to standards including GxP, and data governance.	Published January 2026 (EMA page: first published 14 January 2026; FDA page current as of 14 January 2026). High-level principles, not a guideline or regulation (checked 30 October 2026).	Not addressed. Its scope is AI used to generate or analyse evidence across the drug product life cycle, naming nonclinical, clinical, post-marketing and manufacturing phases; it does not mention discovery (EMA's summary page refers to early research).	Yes. Principle 3 calls for AI to adhere to relevant standards, including Good Practices (GxP).	https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development
Reflection paper on the use of AI in the medicinal product lifecycle (European Medicines Agency)	Principles for AI and machine learning across the lifecycle, from drug discovery to post-authorisation, including data, model development, validation and regulatory interaction.	Final reflection paper, EMA/CHMP/CVMP/83833/2023. Adopted by CHMP on 9 September 2024 and by CVMP on 11 September 2024, published 30 September 2024 (checked 30 October 2026).	Partly. Discovery use can be low regulatory impact if poor performance only affects the developer; if results feed evidence for regulatory review, non-clinical principles apply.	Yes. EMA states it is the responsibility of the clinical trial sponsor, marketing authorisation applicant or holder, or manufacturer to ensure algorithms, models, datasets and pipelines are fit for purpose and in line with EU legislation, GxP standards and EMA guidelines.	https://www.ema.europa.eu/en/use-artificial-intelligence-ai-medicinal-product-lifecycle-scientific-guideline
GAMP Guide: Artificial Intelligence (ISPE)	A risk-based life-cycle approach for AI-enabled computerised systems in GxP areas, covering the regulated company's and suppliers' responsibilities. Used alongside GAMP 5 (Second Edition).	Published industry guide, July 2025, 290 pages. Not a regulation (checked 30 October 2026).	Not its focus: it is written for GxP-regulated systems.	Yes. An industry reference for AI-enabled computerised systems in GxP areas, designed to be used alongside GAMP 5.	https://ispe.org/publications/guidance-documents/gamp-guide-artificial-intelligence
EU GMP Annex 22: Artificial Intelligence (European Commission)	AI models in critical GMP manufacturing applications for medicinal products and active substances. Covers static, deterministic machine-learning models; says dynamic, probabilistic, generative AI and large language models should not be used in critical GMP applications.	Draft. Consultation ran 7 July to 7 October 2025 and is closed; Annex 22 is not yet listed in EudraLex Volume 4 (checked 30 October 2026). The final text may differ from the draft.	No. Its scope is manufacturing.	GMP manufacturing only, and only once adopted. If generative AI or large language models are used in non-critical GMP uses, the draft says qualified, trained staff should always be responsible for their outputs.	https://health.ec.europa.eu/consultations/stakeholders-consultation-eudralex-volume-4-good-manufacturing-practice-guidelines-chapter-4-annex_en
ISO/IEC 42001:2023, AI management system (ISO and IEC)	Requirements for establishing, implementing, maintaining and continually improving an AI management system within an organisation that develops, provides or uses AI.	Published international standard (ISO stage 60.60), Edition 1, December 2023, no amendment listed (ISO's page as archived 6 October 2026; iso.org blocks automated checks, so it could not be re-checked on 30 October 2026).	Yes, as an organisation-wide management system covering any AI use.	Not a GxP requirement. It can sit alongside a pharmaceutical quality system but does not replace computerised-system validation.	https://www.iso.org/standard/42001
AI Risk Management Framework, AI RMF 1.0, NIST AI 100-1 (US NIST)	Voluntary framework built on four functions: GOVERN, MAP, MEASURE and MANAGE. Companion Generative AI Profile, NIST AI 600-1 (July 2024).	Final, released 26 January 2023, intended for voluntary use. NIST states it is being revised as part of the White House AI Action Plan (checked 30 October 2026).	Yes. Voluntary and usable for any AI, including research agents.	Not a GxP requirement. Useful as shared vocabulary for an agent governance programme.	https://www.nist.gov/itl/ai-risk-management-framework

Separately, 21 CFR Part 11 applies to electronic records required by FDA predicate rules or submitted to FDA, and EU GMP Annex 11 to computerised systems used in EU GMP activities, whatever AI-specific guidance says. The EU AI Act is outside this map. Not legal or regulatory advice. Full answer: atlasbiosys.com/answers/ai-agent-governance-standards-pharma-rd.html. Atlas Bio Systems, 2026.