Audit-Trail and Access-Control Checklist for AI Agents in GxP Research

Atlas Bio Systems · Version 1.0 · Published 17 October 2026 · Free to copy and use inside your organization

Short answer. Give every AI agent its own identity, a named human owner and only the access its task needs. Have the platform, not the agent, write a computer-generated, time-stamped entry for every action that creates, changes or deletes a GxP record. Keep a person as the only signer and approver, and review those audit trails under a written procedure. This checklist turns that into 42 items, each citing the source it draws on and marking where that is our reading or a draft text.

How to use this checklist

- Work through one agent or one agentic workflow at a time. Tick an item only when you can show the evidence, not when it is planned.
- GxP rules were written for people and computerised systems, and the clauses below speak of operators, users and individuals. None of them names AI agents. The EU draft Annex 22 (not yet adopted) addresses AI models in GMP manufacturing, covers only static models with deterministic output, and excludes generative AI and large language models from critical GMP use. Applying these clauses to agent actions is our reading, marked “by analogy” or “our reading” where the clause addresses something else.
- The sources come from different regimes. EU GMP Annex 11, the draft Annex 11 and Annex 22, PIC/S PI 041-1 and the FDA data integrity Q&A are GMP (and, for PIC/S, GDP) texts; ICH E6(R3), the EMA computerised systems guideline and the FDA clinical investigations Q&A cover clinical trials; 21 CFR Part 58 and OECD GLP No. 17 cover nonclinical (GLP) studies. Outside its own regime a text is cited as recognised good practice, not as the rule that applies: check the GLP or GCP equivalent. 21 CFR Part 11 and the MHRA guidance span GxP; NIST and OWASP are security frameworks, not GxP rules.
- Items whose only basis is a draft (draft Annex 11, draft Annex 22, FDA's January 2025 draft AI guidance) describe the expected direction, not current requirements. As of 17 October 2026 the EudraLex Volume 4 page still lists Annex 11 at its January 2011 revision and does not list Annex 22.
- Where an item cites 21 CFR 11.10(a), (b), (c) or (e), note that FDA's 2003 Part 11 scope guidance says FDA intends to exercise enforcement discretion over Part 11's validation, audit-trail, record-retention and record-copying requirements; the records must still meet the underlying predicate rules, and Part 11 remains in effect.
- Your quality organisation (the quality unit under GMP, test facility management and the QA unit under GLP, sponsor or CRO QA under GCP) decides which items apply and how. This is a working aid, not legal or regulatory advice, a validation protocol or evidence of compliance.

1. Decide the scope first

☐ **1.** List every agent and every record it can create, modify, delete or transmit.
Basis: Part 11 applies to electronic records kept under FDA record requirements or submitted to FDA (21 CFR 11.1(b)). [1]

☐ **2.** Mark which of those records are required by a predicate rule (GLP, GMP, GCP) or submitted to a regulator, and write the boundary down: exploratory research that never touches a GxP record can be governed more lightly.
Basis: FDA's 2003 Part 11 scope guidance interprets Part 11 narrowly and ties it to predicate-rule records. It says FDA intends to exercise enforcement discretion over Part 11's validation, audit-trail, record-retention and copying requirements (11.10(a), (b), (c), (e)) while those records must still meet predicate rules; Part 11 itself remains in effect. [2]

☐ **3.** For each regulated use, record whether it is critical: a direct impact on patient safety, product quality or data integrity.
Basis: EU GMP Annex 11, section 1 (risk management); draft Annex 22, section 1. [3, 4]

☐ **4.** If the agent runs on a generative or large language model, note where it would sit if the EU draft Annex 22 is adopted. As drafted, it applies only to static models with deterministic output; it says generative AI and large language models should not be used in critical GMP applications, and that in non-critical use qualified personnel stay responsible for checking their outputs (human in the loop). Our reading: an agent that creates, changes or deletes GMP records may itself be a critical application. The draft ties human-in-the-loop for generative models only to non-critical use and does not say whether human approval of each action makes a use non-critical. It defines critical by direct impact on patient safety, product quality or data integrity; it does not list controls around the model, such as human approval, as a factor. Its other human-in-the-loop provisions (sections 3.3 and 10.5) apply to the static, deterministic models within its scope where testing effort has been reduced, adding operator responsibility and record-keeping; for LLMs the draft says only that, in non-critical use, its principles “may be considered where applicable”, and it does not offer them as a route into critical use.
Basis: Draft Annex 22, sections 1, 3.3 and 10.5. It is drafted for GMP manufacturing and would not by itself govern research use. Consultation ran 7 July to 7 October 2025; the text is a draft, not in force. [4, 5, 6]

☐ **5.** Where AI output will support a regulatory submission, write down the model's context of use.
Basis: FDA draft guidance (January 2025) on AI to support regulatory decision-making describes a risk-based credibility assessment tied to the context of use (draft, not in force). [7]

2. Identity and ownership

- ☐ **6.** Each agent has its own unique account. It never runs under a person's login or a shared service account that can change data.
  Basis: FDA data integrity Q&A, question 5 (shared logins); draft Annex 11, 11.1 (shared accounts other than read-only "constitute a violation of data integrity"; draft, not in force). [8, 9]
- ☐ **7.** Each agent has a named human owner who is accountable for what it does, and a written purpose.
  Basis: Annex 11, section 2 and glossary (process owner, system owner); NIST AI RMF, GOVERN 2.1 (documented roles and responsibilities). [3, 10]
- ☐ **8.** Audit-trail entries let a reviewer tell the agent's actions apart from those of the person accountable for it, so records stay attributable in the ALCOA+ sense.
  Basis: PIC/S PI 041-1, 7.5: attributable means it is possible to identify "the individual or computerised system" that performed a recorded task; MHRA GxP data integrity guidance, 3.10 (ALCOA+). [11, 12]
- ☐ **9.** The agent's identity record includes the model, model version, prompts or instructions, tools and configuration in force, so any action can be traced to the exact version that took it.
  Basis: Our reading of Annex 11, section 10 (changes, including configuration, made under a defined procedure); the clause does not specify identity-record contents. [3]

3. Access control

- ☐ **10.** Each agent has an explicit permission list: what it may read, write, call and send, and where. Everything else is denied.
  Basis: NIST SP 800-53 AC-6 (least privilege, including processes acting on behalf of users); draft Annex 11, 11.10 (least privilege); 21 CFR 11.10(d); Annex 11, 12.1. Deny-by-default is our reading. [13, 9, 1, 3]
- ☐ **11.** Authority checks are enforced by the system at the moment of each action, not left to the agent's instructions.
  Basis: 21 CFR 11.10(g) (authority checks so only authorized individuals can alter a record or perform the operation at hand), applied by analogy to agents. [1]
- ☐ **12.** Granting, changing and revoking an agent's access are each recorded.
  Basis: Annex 11, 12.3; ICH E6(R3), 4.3.8(c); OECD GLP No. 17, para. 92. [3, 14, 15]
- ☐ **13.** Agent access is reviewed on a schedule and removed when no longer needed.
  Basis: ICH E6(R3), 4.3.8(b); OECD GLP No. 17, para. 92; draft Annex 11, 11.2 and 11.11 (recurrent reviews; draft). [14, 15, 9]
- ☐ **14.** Agent credentials (API keys, tokens) are kept secret and rotated, and can be disabled at once if compromised; attempts at unauthorized use are detected and reported.
  Basis: Draft Annex 11, 11.4 (authentication kept confidential; draft); by analogy with 21 CFR 11.300(b) to (d), which cover identification codes and passwords. [9, 1]
- ☐ **15.** Content the agent reads (documents, web pages, tool output) is treated as data, never as instructions; the source of any instruction it acts on is checked.
  Basis: OWASP LLM01 (prompt injection); by analogy with 21 CFR 11.10(h) (device, e.g. terminal, checks on the validity of the source of data input or operational instruction). [16, 1]
- ☐ **16.** No agent can verify or approve its own output, or change audit-trail settings.
  Basis: NIST SP 800-53 AC-5 (separation of duties); draft Annex 11, 11.10 (segregation of duties) and 12.3 (draft); OECD GLP No. 17, para. 81 (study personnel should not be authorised to change audit-trail settings). [13, 9, 15]
- ☐ **17.** A named person can suspend any single agent immediately without stopping the rest of the system.
  Basis: NIST AI RMF, MANAGE 2.4 (mechanisms to supersede, disengage or deactivate AI systems). [10]

4. Signatures and approvals stay with people

- ☐ **18.** Agents never apply an electronic signature. Every GxP approval or release, and every review the predicate rule requires to be signed, is signed by a person.
  Basis: 21 CFR 11.100(a): each electronic signature is unique to one individual and not reused or reassigned. [1]
- ☐ **19.** The signed record shows the signer's printed name, the date and time, and the meaning of the signature (review, approval, responsibility, authorship), and the signature is linked to its record.
  Basis: 21 CFR 11.50 and 11.70; Annex 11, section 14. [1, 3]
- ☐ **20.** Actions that change GxP records or send data outside the organization are defined in advance and route to a person for approval.
  Basis: NIST AI RMF, GOVERN 2.1; our practice. As we read draft Annex 22 (see item 4), human approval alone may not make an LLM acceptable in a critical GMP use. [10]

5. What every audit-trail entry records

For each action that creates, modifies or deletes a GxP record. Draft Annex 11 frames these fields for manual user interactions; applying them to agent actions is our reading.

- ☐ **21.** Who: the agent's ID, its role, and the named person accountable for the run (the person who started it or, for a scheduled run, the agent's owner).
  Basis: Annex 11, 12.4; draft Annex 11, 12.2 (the user, including role); 21 CFR 58.130(e) (for GLP, the responsible individual must be identified); FDA data integrity Q&A, question 5. [3, 9, 17, 8]
- ☐ **22.** What: the record and field, the old and new value, and the tool call that made the change.
  Basis: Draft Annex 11, 12.2 (old and new value; draft). [9]
- ☐ **23.** When: a computer-generated time stamp with an unambiguous time zone (for example, UTC).
  Basis: 21 CFR 11.10(e), applied by analogy to agents; ICH E6(R3), 4.2.2(d); FDA clinical investigations Q&A, question 14 (system date and time controls). [1, 14, 18]
- ☐ **24.** Why: the reason for any change or deletion.
  Basis: Annex 11, section 9; 21 CFR 58.130(e) for GLP data. [3, 17]
- ☐ **25.** On what: the input data, documents and prompts used, and the model version.
  Basis: Our reading of what a credibility assessment needs (FDA draft AI guidance); not a stated Part 11 field. [7]
- ☐ **26.** Outcome: what the agent produced, and who accepted or rejected it, and when.
  Basis: 21 CFR 11.10(e) (operator entries and actions), applied by analogy to agents; PIC/S PI 041-1, 7.5 (attributable). [1, 11]
- ☐ **27.** The entry is written by the platform or gateway at the time of the action, not composed afterwards by the agent itself.
  Basis: 21 CFR 11.10(e) (computer-generated, independently recorded), applied by analogy; draft Annex 11, 12.2 (recorded at the time of events; draft). [1, 9]

6. Protect the audit trail

- ☐ **28.** Audit trails are on at all times and cannot be disabled or edited by the agent, its owner or ordinary users; any change to settings or system time creates its own entry.
  Basis: ICH E6(R3), 4.2.2(b); EMA computerised systems guideline (clinical trials), 6.2.1; draft Annex 11, 12.3 (draft). [14, 19, 9]
- ☐ **29.** A change never obscures what was recorded before.
  Basis: 21 CFR 11.10(e); 21 CFR 58.130(e); OECD GLP No. 17, para. 80. [1, 17, 15]
- ☐ **30.** Audit data and logging tools are protected from unauthorized access, change and deletion, with an alert when unauthorized access, change or deletion is detected.
  Basis: NIST SP 800-53 AU-9. [13]

7. Review the audit trail

- ☐ **31.** A written procedure states who reviews agent audit trails, what they review and when.
  Basis: FDA data integrity Q&A, question 7; EMA computerised systems guideline, 6.2.2; OECD GLP No. 17, para. 82; draft Annex 11, 12.5. [8, 19, 15, 9]
- ☐ **32.** Review frequency follows the predicate rule where one sets it; otherwise it is set by documented risk assessment.
  Basis: FDA data integrity Q&A, question 8. [8]
- ☐ **33.** Reviewers were not directly involved in the activity they review.
  Basis: Draft Annex 11, 12.6 (draft, not in force). [9]
- ☐ **34.** Audit trails can be exported in human-readable and electronic form, and searched and sorted by who, what, when and why.
  Basis: 21 CFR 11.10(b); Annex 11, section 9; EMA computerised systems guideline, 6.2.1; FDA clinical investigations Q&A, question 12; draft Annex 11, 12.4. [1, 3, 19, 18, 9]
- ☐ **35.** Unexpected findings from a review are investigated and the outcome recorded.
  Basis: Draft Annex 11, 12.5 (draft, not in force). [9]

8. Retention, incidents and change

- ☐ **36.** Agent audit trails are kept at least as long as the records they cover, and stay available for inspection.
  Basis: 21 CFR 11.10(c) and (e); the retention period itself is set by the predicate rule (for example 21 CFR 58.195 for GLP, 211.180 for GMP, 312.57(c) for sponsors and 312.62(c) for clinical investigators) and, in the EU, by the applicable GxP rule; FDA clinical investigations Q&A, questions 5 and 12 (retain records, metadata and audit trails). [1, 17, 20, 21, 18]
- ☐ **37.** Archived audit trails remain readable after the model, platform or vendor changes, and this is tested.
  Basis: Annex 11, section 17. [3]
- ☐ **38.** Every incident involving an agent is reported and assessed; for critical incidents the root cause feeds corrective and preventive action. Logs are preserved before the investigation starts (our practice, not stated in Annex 11).
  Basis: Annex 11, section 13. [3]

☐ **39.** The agentic system is validated for its intended use, with effort scaled to a documented risk assessment.
Basis: 21 CFR 11.10(a); Annex 11, sections 1 and 4. [1, 3]

☐ **40.** A change of model, model version, prompt, tool or permission is a controlled change under a defined procedure.
Basis: Annex 11, section 10. [3]

☐ **41.** Written agreements with AI model and platform providers state each party's responsibilities.
Basis: Annex 11, 3.1 (formal agreements with third parties, including for data processing); draft Annex 11, 2.6 (the regulated user stays fully responsible; draft). [3, 9]

☐ **42.** People who own, configure or review agents are trained for that task.
Basis: 21 CFR 11.10(i). [1]

Sign-off

Agent or workflow reviewed:	
Reviewed by (name, role):	
Date:	
Next review due:	

Sources

[1] eCFR: 21 CFR Part 11, Electronic Records; Electronic Signatures. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11

[2] FDA guidance: Part 11, Electronic Records; Electronic Signatures - Scope and Application (2003). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/part-11-electronic-records-electronic-signatures-scope-and-application

[3] European Commission: EU GMP Annex 11, Computerised Systems (2011). https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf

[4] European Commission: draft Annex 22, Artificial Intelligence (consultation draft, 2025). https://health.ec.europa.eu/document/download/5f38a92d-bb8e-4264-8898-ea076e926db6_en?filename=mp_vol4_chap4_annex22_consultation_guideline_en.pdf

[5] European Commission: stakeholder consultation on EudraLex Volume 4 Chapter 4, Annex 11 and new Annex 22. https://health.ec.europa.eu/consultations/stakeholders-consultation-eudralex-volume-4-good-manufacturing-practice-guidelines-chapter-4-annex_en

[6] European Commission: EudraLex Volume 4, EU Good Manufacturing Practice guidelines. https://health.ec.europa.eu/medicinal-products/eudralex/eudralex-volume-4_en

[7] FDA draft guidance: Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products (January 2025). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

[8] FDA guidance: Data Integrity and Compliance With Drug CGMP - Questions and Answers (2018). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/data-integrity-and-compliance-drug-cgmp-questions-and-answers

[9] European Commission: draft revised Annex 11, Computerised Systems (consultation draft, 2025). https://health.ec.europa.eu/document/download/40231f18-e564-4043-94de-c031f813d38b_en?filename=mp_vol4_chap4_annex11_consultation_guideline_en.pdf

[10] NIST AI 100-1: Artificial Intelligence Risk Management Framework (AI RMF 1.0). https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

[11] PIC/S PI 041-1: Good Practices for Data Management and Integrity in Regulated GMP/GDP Environments (2021). https://picscheme.org/docview/4234

[12] MHRA: GXP Data Integrity Guidance and Definitions (March 2018). https://www.gov.uk/government/publications/guidance-on-gxp-data-integrity

[13] NIST SP 800-53 Rev. 5: Security and Privacy Controls for Information Systems and Organizations. https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final

[14] ICH E6(R3): Guideline for Good Clinical Practice (final, January 2025). https://database.ich.org/sites/default/files/ICH_E6%28R3%29_Step4_FinalGuideline_2025_0106.pdf

[15] OECD Series on Principles of GLP No. 17: Application of GLP Principles to Computerised Systems, ENV/JM/MONO(2016)13, reissued 9 November 2022 (copy hosted by Germany's BfR). https://www.bfr.bund.de/cm/349/No17_Application-of-GLP-Principles-to-Computerised-Systems_2022.pdf

[16] OWASP GenAI Security Project: LLM01 Prompt Injection. https://genai.owasp.org/llmrisk/llm01-prompt-injection/

[17] eCFR: 21 CFR Part 58, Good Laboratory Practice for Nonclinical Laboratory Studies. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-58

[18] FDA guidance: Electronic Systems, Electronic Records, and Electronic Signatures in Clinical Investigations - Questions and Answers (October 2024). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/electronic-systems-electronic-records-and-electronic-signatures-clinical-investigations-questions

[19] EMA: Guideline on computerised systems and electronic data in clinical trials (2023). https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-computerised-systems-and-electronic-data-clinical-trials_en.pdf

[20] eCFR: 21 CFR Part 211, Current Good Manufacturing Practice for Finished Pharmaceuticals. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-C/part-211

[21] eCFR: 21 CFR Part 312, Investigational New Drug Application. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-D/part-312

About

Atlas Bio Systems builds agentic AI for life-science research. This checklist is published as a free working aid. It records the clauses as read on 17 October 2026; draft texts may change before adoption. The plain-language answer this checklist expands: Audit Trails for AI Agents in Regulated Research. Questions: contact Atlas Bio.
Latest version: https://atlasbiosys.com/downloads/ai-agent-gxp-audit-trail-access-control-checklist.pdf