PROCESS · GOVERNANCE

Write a One-Page AI Policy Your SOPs Can Reference

An OmniTek Consulting implementation guide for VPs and Directors in pharma, biotech, and medical device companies.

READING TIME
About 10 minutes

WHO IT'S FOR
VPs and Directors

GOES WITH
The AI Maturity Quick Check

> A one-page AI policy tells your team which AI tools they can use, with which data, for what work, and who checks the output. Written well and referenced from your SOPs, it speeds adoption up, because people stop guessing what's allowed. You can have a working version in 30 days.

SCOPE This guide covers AI used in how your company works, such as drafting, analysis, and operations. It doesn't cover AI that is part of a product you sell, such as AI-enabled device software or software as a medical device. That has its own regulatory path; check it with your Regulatory team.

Why it matters

Most life sciences teams are already using AI, whether or not there's a policy. Without one, three things tend to happen:

- **People hold back.** Careful employees avoid AI entirely because they don't know what's permitted, so the department falls behind.
- **People improvise.** Others paste trial data, draft deviation reports, or summarize regulatory correspondence in personal accounts nobody approved.
- **Quality gets surprised.** The first time Quality hears about an AI use is during an audit or an inspection, which is the worst time to have that conversation.

A long enterprise policy rarely fixes this, because no one reads it. One page that your team can actually remember, tied into the procedures they already follow, does.

What the policy covers

A usable policy answers seven questions. Each needs only a few lines.

Part	The question it answers	Example line
1. Scope	Who and what does this cover?	"Applies to all [Department] staff and contractors using AI tools for [Company] work."
2. Approved tools	Which tools can I use, and how do I request a new one?	"Use only the tools on the approved list. Request new tools through IT and Quality."
3. Data rules	What information can go into which tool?	"Never enter patient, trial participant, or unpublished study data into a tool not approved for it."
4. Prohibited uses	What is off limits, full stop?	"AI may not make final decisions on product release, batch disposition, deviation or CAPA closure, approval of SOPs or master records, supplier approval, patient eligibility or unblinding, patient safety and adverse event assessment, or regulatory submissions."
5. Human review	Who checks AI output, and when?	"A named person reviews and is accountable for any AI output used in a deliverable."
6. GxP uses	What extra steps apply in regulated work?	"Any AI use in a GxP process is classified with Quality before use and follows the validation SOP."
7. Ownership	Who owns this, how are people trained, and how often is it updated?	"Owned by [Name]. Training is recorded. Reviewed every six months in year one, then yearly, and after a new tool, a rule or classification change, or an AI-related deviation."

Parts 3 and 6 are where life sciences policies differ most from generic ones. Get those right with Quality, and the rest is straightforward.

A one-page template

Copy this, replace the bracketed text, and cut anything that doesn't apply to your department.

TEMPLATE · COPY AND ADAPT

[Department] AI Use Policy

Owner: [Name, title] | Effective: [date] | Next review: [date] | Referenced in: [SOP numbers]

1. Scope

Applies to all [Department] staff and contractors using AI tools for [Company] work.

2. Approved tools

Use only the tools on the approved list at [location]. Today that includes [Tool A] for [general drafting and summarizing]. To request a new tool, submit [form or ticket] to IT and Quality. Personal AI accounts may not be used for company work.

3. Data rules

- Public information: any approved tool.
- Internal information: approved enterprise tools only, with contract terms that stop vendors training on company data.
- Confidential information: [approved enterprise tools only].
- Unpublished study results: aggregated results with no participant data or partner restrictions are Confidential. Participant-level data is Personal and patient. Results under a partner or licensing confidentiality clause, and topline results not yet publicly disclosed, stay Confidential, with an extra rule: they may go only into [named tools], or not into AI tools at all.
- Personal and patient data, including trial participant data: [named tools only, or not permitted].
- GxP records: only within validated systems.

4. Prohibited uses

AI may not make final decisions on product release, batch disposition, deviation or CAPA closure, approval of SOPs or master records, supplier approval, patient eligibility or unblinding, patient safety and adverse event assessment, or regulatory submissions. AI may not delete or alter raw data or audit trails, or override equipment alarms. AI may not generate, modify, or finalize GxP records outside a validated system. AI drafts are allowed when a qualified person reviews them and enters the record through the validated system.

5. Human review

The person who uses AI output in a deliverable is accountable for it. Check facts, citations, and numbers against source documents before use. [Deliverables X and Y] need a second reviewer.

6. GxP uses

Before any AI use in a GxP process, classify it with Quality using [classification SOP]. GxP uses follow [validation SOP] and change control. Record the classification.

7. Training and questions

All staff complete [training] before using AI tools for work. Training is recorded in [system]. Questions or concerns go to [owner or channel]. This policy is reviewed every six months in year one, then yearly, and after a new tool, a rule or classification change, or an AI-related deviation.

Your first 30 days

Five steps over four weeks. Each step takes a few hours of your time, not a project team.

1. **Week 1: Find out what's already happening.** Send a short anonymous survey asking which AI tools people use, for what tasks, and with what kinds of data. Promise no one gets in trouble for answering honestly. The answers tell you which rules matter most.

2. **Week 1: Agree on the basics with IT and Quality.** In one meeting, agree which tools are approved, which data classes can go into each, and who in Quality will classify GxP uses. These three answers fill most of the template.
3. **Week 2: Draft the page.** Fill in the template above. Keep it to one page; move detail into existing SOPs rather than into the policy.
4. **Week 3: Review and connect.** Have Quality, Legal or Privacy, and IT review the draft. Route it through document control so the relevant SOPs reference it (see the next section).
5. **Week 4: Roll it out.** Walk the team through it in a 30–minute session using real examples from the Week 1 survey. Record the training, publish the policy where SOPs live, and put the review dates (every six months in year one, then yearly) in the owner's calendar.

By day 30 you have a policy people know, an approved tool list, a GxP classification route, and a training record. That covers the governance question in the Quick Check at the Operationalizing level.

Connecting it to your SOPs and quality system

The policy works when it is a controlled document that your existing SOPs point to, not a standalone memo. Give it a document number and version, put it under change control, and add a one-line reference in the SOPs where AI use actually shows up.

Existing procedure	Reference to add
Computerized systems or validation SOP	AI tools used in GxP processes are classified per the AI policy before validation planning.
Data integrity SOP	Data rules for AI tools are defined in the AI policy, section 3.
Supplier qualification SOP	AI vendors are evaluated against the AI policy's approved-tool criteria.
Training SOP	The AI policy is part of the required curriculum for [Department] roles.
Deviation and CAPA SOP	Errors traced to AI output are logged, investigated for root cause, and fixed. In GxP processes they are handled as deviations under this procedure and noted as AI-related; in other processes they go through the normal incident or issue process.
Document control SOP	The AI policy is reviewed every six months in its first year, then yearly, and whenever a new AI tool is approved, a GxP classification or relevant rule changes (for example, Annex 22 becoming final), or an AI-related deviation occurs.

This keeps the policy short. Detail lives in procedures your auditors already know, and a change to one tool doesn't mean rewriting the policy.

Common mistakes

- **Banning everything.** A policy that says "no AI" doesn't stop use; it pushes it into personal accounts where you can't see it. Approve at least one secure tool at the same time you publish the policy.

- **Treating every use as GxP.** Drafting a training outline and assisting a batch record review are not the same risk. Without a classification step, low-risk work gets buried in validation and nothing moves.
- **Writing a 20-page policy.** Long policies go unread. Keep the policy to one page and let SOPs carry the detail.
- **Naming a tool instead of a process.** Tools change every few months. Write "approved tools listed in [location]" so a new tool means updating a list, not reissuing the policy.
- **Skipping the training record.** In a regulated department, a policy without recorded training is hard to defend in an inspection.
- **No owner.** If no one owns the review date, the policy is out of date within six months.

Regulatory touchpoints

No regulator requires a one-page AI policy, but several existing and coming rules expect what it contains. Status is as of October 2026; several of these are still moving.

Rule	Status	Where it shows up in your policy
21 CFR Part 11 (electronic records)	In force	Part 3: GxP records stay in validated systems with audit trails and access controls
FDA Computer Software Assurance guidance	Final. The current version (February 2026) is titled "Computer Software Assurance for Production and Quality Management System Software" and refers to FDA's Quality Management System regulation (21 CFR Part 820); written for device production and quality system software, widely used as a model for risk-based assurance	Part 6: classify first, then test in proportion to risk
FDA draft guidance on AI supporting regulatory decisions for drugs and biologics	Draft since January 2025, still draft as of October 2026 (status)	Parts 4 and 5: if AI produces data or information that supports an FDA decision on safety, effectiveness, or quality, assess its credibility for that use, scaled to risk. AI used for drug discovery or operational efficiency that doesn't affect safety, quality, or study reliability is out of scope. Named human review is this policy's own control, not an FDA term
EU GMP Annex 22 (AI), revised Annex 11, and Chapter 4	Drafts published for consultation in July 2025; not yet final (status)	Parts 4 and 6: the draft limits critical GMP uses to fixed, repeatable models, keeps generative AI out of critical functions, and requires human oversight
EU AI Act, Article 4 (AI literacy)	Applies to providers and deployers of AI systems in the EU (Regulation (EU) 2024/1689). As amended in July 2026 by Regulation (EU) 2026/1744, it requires them to take measures to support AI literacy; it doesn't require a set level for any individual.	Part 7: recorded training shows the measures you took

This guide is general information, not legal or regulatory advice. Have your Quality and Regulatory leads confirm how these apply to your products and markets.



How OmniTek can help

A policy is the first governance step. The full AI RMaturityReport Card shows where your department stands across all 13 maturity dimensions, including GxP validation, and comes with the **AI policy kit**: an editable version of this template, ready-to-use SOP reference language, and a rollout email for your team.

To talk through your policy or your Quick Check results, book time with OmniTek.

Book time with OmniTek

SOURCES

- FDA: Computer Software Assurance for Production and Quality Management System Software (final guidance, February 2026)
- FDA draft guidance: Considerations for the Use of AI to Support Regulatory Decision-Making for Drug and Biological Products
- FDA guidance page: Considerations for the Use of AI to Support Regulatory Decision-Making for Drug and Biological Products (status: draft)
- European Commission: Stakeholders' consultation on EudraLex Volume 4, Chapter 4, Annex 11 and new Annex 22 (July–October 2025)
- EUR-Lex: Regulation (EU) 2026/1744 (Digital Omnibus on AI), adopted 8 July 2026; in force 27 July 2026
- eCFR: 21 CFR Part 11

This guide is general information for leaders in regulated life sciences functions, not legal or regulatory advice. Confirm requirements with your Quality, Regulatory, and Legal teams. © 2026 OmniTek Consulting.