AI vendor data-handling questionnaire for life-science R&D

Published by Atlas Bio Systems (atlasbiosys.com) · Version 2026-10-19 · https://atlasbiosys.com/secure-agentic-ai-life-sciences

Questions to put to any AI vendor before research data reaches its product: training use, retention, sub-processors, tenancy and isolation, audit logs, agent permissions, GxP and 21 CFR Part 11 support, incident notice, and exit and deletion. Each question says why it matters and points to the closest provision in a public framework. The Excel version has columns for the vendor's answers and your assessment.

How to use it

• Send the Questionnaire sheet to the vendor before you share any proprietary or partner data. Ask for the contract clause or document that backs each answer, not only a yes or no.
• Area G applies only where the AI touches GxP-regulated work. Mark it N/A otherwise. In other areas, references marked [GxP scope only] are given as good practice: EU GMP Annex 11 binds only for computerised systems used in GMP-regulated activities, and 21 CFR Part 11 only for records required by FDA predicate rules or submitted to FDA. Your quality function decides whether a use is in scope.
• Annex 11 is written for GMP. For clinical trials, the EU counterpart is EMA's Guideline on computerised systems and electronic data in clinical trials (2023); for non-clinical safety studies under GLP, see OECD GLP Advisory Document No. 17 on computerised systems. In the US, Part 11 applies to GLP, GCP and GMP predicate-rule records alike. The European Commission's draft GMP Annex 22 on artificial intelligence (a consultation draft, not in force) does not apply to generative AI or large language models and says they should not be used in critical GMP applications, those with direct impact on patient safety, product quality or data integrity; where they are used in non-critical GMP applications, qualified personnel should be responsible for checking that outputs are suitable.
• Use the Assessment column to record Meets, Partly meets, Does not meet, N/A or Not answered, and note what you would need to accept the gap.
• Agree with your quality, legal and security teams which answers are deal-breakers before the replies arrive.

General guidance, not legal or regulatory advice. This questionnaire does not validate any system or certify any vendor. It is generic and was built from the public frameworks listed under Sources; each framework reference is the closest provision, not a statement that the provision applies to your use. Decisions on GxP scope and validation stay with your quality function.

A. Training use

ID	Question	Why it matters	Framework reference
A1	Will our prompts, outputs, uploaded files or retrieved content be used to train, fine-tune or otherwise improve any model, yours or a third party's? Is that off by default on the plan we are buying, and which contract clause says so?	Data that has been used for training is hard to remove reliably afterwards. OWASP lists unclear terms and privacy policies that let a model operator train on an application's sensitive data as a supply-chain risk.	OWASP LLM03:2025 (unclear T&Cs and data privacy policies of model operators); OWASP LLM02:2025 (let users opt out of having their data included in training); NIST AI RMF GOVERN 6.1 (policies for AI risks from third-party entities)
A2	Apart from training, do your staff, contractors or sub-processors use our data for anything else, such as human review, quality evaluation, abuse monitoring or product analytics? Under what conditions, and can we opt out?	Training is not the only secondary use. Each extra use is another person or system that sees unpublished research data.	NIST AI 600-1 Data Privacy risk (leakage and unauthorized use or disclosure of sensitive data); NIST AI RMF MEASURE 2.10 (privacy risk examined and documented)
A3	If your data-use terms change, how far in advance do you notify us, and can we leave before the new terms apply to our data?	A good answer today can change at the next terms update. OWASP recommends regularly reviewing suppliers for changes in their security posture or terms.	OWASP LLM03:2025 (regularly review and audit supplier security, access and T&Cs); NIST AI RMF MANAGE 3.1 (third-party AI risks regularly monitored)
A4	Could our proprietary content, such as sequences, structures, assay results or unpublished hypotheses, appear in answers given to other customers?	For R&D teams the core exposure is trade secrets. NIST names eased exposure of trade secrets as a generative-AI intellectual-property risk.	NIST AI 600-1 Intellectual Property risk (eased exposure of trade secrets); OWASP LLM02:2025 (sensitive information disclosed through output)

B. Retention

ID	Question	Why it matters	Framework reference
B1	How long do you keep each of the following: prompts, outputs, uploaded files, embeddings or vector indexes, caches, logs and backups? Can we set shorter retention?	Each of these is a copy of your data. A short retention period for chat history means little if embeddings or backups are kept indefinitely.	OWASP LLM02:2025 (clear policies on data retention, usage and deletion); OWASP LLM08:2025 (embedding inversion can recover source information)
B2	In which countries or regions is our data stored and processed, including backups, logs and any sub-processor? Can we restrict it?	Partner agreements and data-protection obligations can depend on where data sits. NIST asks organizations to map legal risks of third-party components.	NIST AI RMF MAP 4.1 (legal risks of AI components, including third-party data or software, mapped and documented)
B3	For records we must keep under GxP rules, can they be retrieved accurately and readably for the whole retention period, including after a model, software or platform change? How is restoration from backup tested?	A regulated record has to stay retrievable for years, not just until the next product release.	21 CFR 11.10(c) (protection of records for accurate and ready retrieval throughout the retention period); FDA 2003 Part 11 scope guidance (enforcement discretion on record retention; predicate-rule retention requirements still apply); EU GMP Annex 11 sections 7.1, 7.2 and 17 (data storage, back-ups, archiving) [GxP scope only]

C. Sub-processors and model providers

ID	Question	Why it matters	Framework reference
C1	List every sub-processor and third-party model provider that may receive our data, what each receives and where it is located.	Your contract is with one vendor, but your data may reach several. NIST lists untraceable integration of upstream third-party components and improper supplier vetting as a generative-AI risk.	NIST AI 600-1 Value Chain and Component Integration risk; NIST AI RMF GOVERN 6.1; EU GMP Annex 11 section 3.1 (formal agreements with third parties used for data processing, with clear responsibilities) [GxP scope only]
C2	How do you notify us before adding or changing a sub-processor or underlying model provider, and can we object or leave?	A change of model provider can change where your data goes and under which terms, without any change to the product you see.	NIST AI RMF MANAGE 3.1 (third-party AI risks regularly monitored, controls applied and documented); OWASP LLM03:2025 (regularly review supplier security posture and T&Cs)
C3	Which pre-trained, open-source or third-party models does the product use, which versions, and how do you monitor and update them?	Model components carry their own risks, including poisoned or tampered models. You need to know what is inside the product you are relying on.	NIST AI RMF MANAGE 3.2 (pre-trained models monitored as part of regular monitoring and maintenance); OWASP LLM03:2025 (supply-chain integrity of models)

D. Tenancy and isolation

ID	Question	Why it matters	Framework reference
D1	Is our data separated from other customers' data at every layer: storage, vector database, caches, fine-tuned models and logs? Logically or physically?	OWASP warns that in multi-tenant vector databases, embeddings from one group can be retrieved in response to another group's queries, leaking business information.	OWASP LLM08:2025 (context leakage in multi-tenant environments)
D2	Can the service run in our own cloud tenancy or a dedicated instance? If so, which components, such as model inference, telemetry or support access, still run outside it?	Deployment options decide what leaves your boundary. Ask for a data-flow diagram, not just a label. For GxP records, whether the service is an open or closed system is for your quality function to decide.	NIST AI RMF MAP 4.1 (risks of components, including third-party software, mapped); 21 CFR 11.30 (controls for open systems, where system access is not controlled by those responsible for the records' content; definition in 11.3) [GxP scope only]
D3	Does retrieval honour the permissions of our source systems, so a user cannot see content through the AI that they could not open directly?	An AI assistant connected to shared drives can become a way around access controls. OWASP recommends permission-aware vector stores.	OWASP LLM08:2025 (permission and access control for vector stores); 21 CFR 11.10(d) (limiting system access to authorized individuals) [GxP scope only]
D4	Who at your company, or at a sub-processor, can access our data or environment, under what approval, and is that access logged and shown to us?	Vendor support access is another route to your data. Annex 11 expects physical or logical controls restricting system access to authorized persons.	EU GMP Annex 11 section 12.1 (access restricted to authorised persons) and section 3.1 (third parties that maintain systems, e.g. via remote access) [GxP scope only]

E. Audit logs

ID	Question	Why it matters	Framework reference
E1	For each AI interaction and agent action, what is logged: user, time, input, sources retrieved, tools called, output, model version and approvals? Can we export the logs in a usable format?	When something goes wrong, you need to rebuild what the system saw, did and who approved it. OWASP recommends logging and monitoring what LLM extensions and downstream systems do.	OWASP LLM06:2025 (log and monitor the activity of LLM extensions and downstream systems); NIST AI RMF MEASURE 2.8 (transparency and accountability risks examined); EU GMP Annex 11 section 9 (risk-based audit trail of GMP-relevant changes and deletions, with the reason documented; covers GMP data, not every AI interaction) [GxP scope only]
E2	Are logs time-stamped, protected against change and kept as long as we need? When a record changes, is the earlier value preserved?	A log that can be edited, or that overwrites earlier values, is weak evidence. Part 11 requires that record changes not obscure previously recorded information.	21 CFR 11.10(e) (secure, computer-generated, time-stamped audit trails; changes shall not obscure previous information); FDA 2003 Part 11 scope guidance (enforcement discretion on audit trails; predicate rules still apply); EU GMP Annex 11 sections 9 and 12.4 [GxP scope only]
E3	Are the model version, system prompt and configuration recorded with each output, and are changes to them made under change control that we can see?	The same question can get a different answer after a silent model or prompt change. Without the version you cannot reproduce or defend an output.	NIST AI RMF MANAGE 4.1 (post-deployment monitoring, including change management); EU GMP Annex 11 section 10 (change and configuration management) [GxP scope only]

F. Agent permissions

ID	Question	Why it matters	Framework reference
F1	Which systems, tools and data can an agent reach, and are its permissions limited to the minimum each task needs? Can we set those limits ourselves?	An agent that can read, call tools and send data is a privileged identity. OWASP lists excessive functionality, permissions and autonomy as the causes of excessive agency.	OWASP LLM06:2025 (minimize extension permissions); 21 CFR 11.10(g) (authority checks) [GxP scope only]
F2	Do agent actions run with the permissions of the person they act for, or under a shared service account? Does each logged action record both the person the agent acted for and the agent itself (name and version), so the record never shows a person as having made an entry the agent made?	A shared account hides who did what and can give every user the agent's full access. Running as the person is safer for access, but if the log then names the person as the author of the agent's entry, the record is misattributed.	OWASP LLM06:2025 (execute extensions in the user's context with minimum privileges); EU GMP Annex 11 section 12.4 (record the identity of operators entering, changing, confirming or deleting data) [GxP scope only]
F3	Which agent actions need a person's approval before they take effect, and can we change that list? How is the approval recorded?	OWASP recommends human-in-the-loop approval of high-impact actions before they are taken. You decide what counts as high-impact in your workflows.	OWASP LLM06:2025 (require user approval for high-impact actions); NIST AI RMF MAP 3.5 (processes for human oversight defined, assessed and documented)
F4	How does the product defend against instructions hidden in documents, web pages, emails or tool results that an agent reads?	Prompt injection is first on OWASP's LLM list in both the 2025 and 2026 editions, and agent goal hijack (ASI01) is first on its agentic list. Agents that read outside content are exposed by design.	OWASP LLM01:2025 (prompt injection, including indirect injection from external content); OWASP Top 10 for Agentic Applications for 2026 (agent goal hijack, ASI01)
F5	Can we pause or switch off an agent, or all agents, immediately? Who can do it, and what happens to tasks in progress?	You need a fast way to stop a system that is behaving outside its intended use.	NIST AI RMF MANAGE 2.4 (mechanisms to supersede, disengage or deactivate AI systems); NIST AI 600-1 GV-1.7-001 (protocols so generative AI systems can be deactivated)

G. GxP and 21 CFR Part 11 support

ID	Question	Why it matters	Framework reference
G1	If we use the product in GxP-regulated work, what supplier documentation do you provide to support our validation (requirements, test evidence, release notes, known issues), and how far in advance do you tell us about changes, including model or prompt changes, that could affect the validated state of our system?	The regulated user validates the system for its intended use and should be able to justify that validation based on risk. Annex 11 expects you to assess the supplier and review its documentation. For a vendor-hosted AI service much of that evidence sits with the vendor, but the vendor cannot validate on your behalf. Decisions on GxP scope and validation stay with your quality function.	EU GMP Annex 11 sections 3.3 and 4.5 (regulated user reviews supplier documentation and assesses the supplier), section 4 (validation), section 10 (change and configuration management) and section 11 (periodic evaluation that systems remain in a valid state); 21 CFR 11.10(a) (validation of systems); FDA 2003 Part 11 scope guidance (enforcement discretion on Part 11 validation; predicate rules still apply)
G2	Will you share quality-system and audit information, accept a supplier audit where our risk assessment calls for one, and make that information available to inspectors on request?	Annex 11 bases the need for a supplier audit on risk and expects supplier quality and audit information to be available to inspectors.	EU GMP Annex 11 sections 3.2 and 3.4 (supplier assessment and audit; information available to inspectors)
G3	Where records are signed in or through the product, how is each electronic signature linked to its record and to one identified person, and does it show the signer's name, the date and time, and the meaning (for example review or approval)? Can an agent apply, reuse or bypass a signature, including by acting inside a session in which a person has already signed once?	Under Part 11 an electronic signature is executed, adopted or authorized by an individual, must not be reused by or reassigned to anyone else, and must be linked to its record so it cannot be excised, copied or transferred to falsify a record. An agent should not hold or apply one: it can prepare and route a record, but the person signs. Part 11 lets later signings in one continuous session use at least one component that only that individual can execute, so an agent acting inside that session must not be able to trigger them.	21 CFR 11.3 (electronic signature executed, adopted or authorized by an individual); 11.50 (signature manifestations: printed name, date and time, meaning); 11.70 (signature/record linking); 11.100(a) (unique to one individual; not reused or reassigned); 11.200(a) (signature components; signings in one continuous session; use only by the genuine owner); EU GMP Annex 11 section 14 (electronic signature)
G4	If AI output could support a regulatory decision on safety, effectiveness or quality, can you help document the model's context of use and the evidence for its credibility?	FDA's January 2025 draft guidance ties model credibility to a defined context of use, assessed by risk. It is a draft, not final. It does not address AI used in drug discovery, or AI used for operational efficiencies (such as internal workflows or drafting a submission) that do not affect patient safety, drug quality or the reliability of nonclinical or clinical study results.	FDA draft guidance, January 2025 (risk-based credibility assessment tied to context of use; draft, not final)

H. Incident notice

ID	Question	Why it matters	Framework reference
H1	What do you treat as an incident (for example data exposure, unauthorized access, harmful agent action, serious model error), and within what time will you notify us, in writing and to whom?	Annex 11 expects all incidents, not only system failures and data errors, to be reported and assessed. Put the notice period in the contract.	NIST AI RMF MANAGE 4.3 (incidents and errors communicated to relevant AI actors) and GOVERN 6.2 (contingency processes for third-party incidents); EU GMP Annex 11 section 13 (incident management) [GxP scope only]
H2	After an incident affecting us, will you share the root cause and the corrective and preventive actions?	Notice alone does not tell you whether the problem is fixed or whether your own records are affected.	NIST AI 600-1 MG-4.3-001 (after-action assessments of incidents); EU GMP Annex 11 section 13 (root cause of a critical incident as the basis of corrective and preventive actions) [GxP scope only]
H3	Do you track and tell customers about errors, near misses and model problems, not only security breaches?	In research, a wrong output that someone relies on can matter as much as a breach.	NIST AI 600-1 MG-4.3-002 (record and track reported errors, near-misses and negative impacts); NIST AI RMF MANAGE 4.3

I. Exit and deletion

ID	Question	Why it matters	Framework reference
I1	When the contract ends, can we export all our data, logs, prompts, configurations and agent definitions in a usable format, and how long do we have to do it?	If you cannot take your records and audit trail with you, leaving becomes a compliance problem as well as a commercial one.	NIST AI RMF GOVERN 1.7 (decommissioning and phasing out AI systems safely); EU GMP Annex 11 section 17 (archived data stays accessible, readable and retrievable through system changes) [GxP scope only]
I2	How and when do you delete our data from production, backups, logs, embeddings and sub-processors, and will you confirm deletion in writing?	Deletion that skips backups, embeddings or sub-processors leaves copies behind. NIST lists data leakage after decommissioning as a factor to consider.	NIST AI 600-1 GV-1.7-002 (data retention, data security and leakage after decommissioning); OWASP LLM02:2025 (policies on retention, usage and deletion)
I3	At exit, what happens to any model, adapter or index built from our data: is it deleted, handed over, or kept?	A model fine-tuned on your data, or an index built from it, can carry that data's value with it after you leave.	NIST AI 600-1 GV-1.7-002 (decommissioning factors); NIST AI RMF GOVERN 6.1 (third-party intellectual property and other rights)

Sources

NIST AI RMF: NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST, January 2023, voluntary framework. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

NIST AI 600-1: NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST, July 2024, companion to the AI RMF; actions are suggested, not required. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

OWASP LLM01:2025: OWASP Top 10 for LLM Applications 2025: LLM01 Prompt Injection. OWASP GenAI Security Project, 2025 edition (LLM01 in the 2026 edition). https://genai.owasp.org/llmrisk/llm01-prompt-injection/

OWASP LLM02:2025: OWASP Top 10 for LLM Applications 2025: LLM02 Sensitive Information Disclosure. OWASP GenAI Security Project, 2025 edition (LLM02 in the 2026 edition). https://genai.owasp.org/llmrisk/llm022025-sensitive-information-disclosure/

OWASP LLM03:2025: OWASP Top 10 for LLM Applications 2025: LLM03 Supply Chain. OWASP GenAI Security Project, 2025 edition (LLM04 in the 2026 edition). https://genai.owasp.org/llmrisk/llm032025-supply-chain/

OWASP LLM06:2025: OWASP Top 10 for LLM Applications 2025: LLM06 Excessive Agency. OWASP GenAI Security Project, 2025 edition (LLM03 in the 2026 edition). https://genai.owasp.org/llmrisk/llm062025-excessive-agency/

OWASP LLM08:2025: OWASP Top 10 for LLM Applications 2025: LLM08 Vector and Embedding Weaknesses. OWASP GenAI Security Project, 2025 edition (LLM09 in the 2026 edition). https://genai.owasp.org/llmrisk/llm082025-vector-and-embedding-weaknesses/

OWASP Agentic Top 10: OWASP Top 10 for Agentic Applications for 2026. OWASP GenAI Security Project, December 2025. https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/

EU GMP Annex 11: EudraLex Volume 4, Annex 11: Computerised Systems. European Commission, in operation since 30 June 2011; GMP guidance. Section numbers refer to the 2011 text in force. https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf

21 CFR Part 11: 21 CFR Part 11, Electronic Records; Electronic Signatures. FDA regulation, in force; applies to records required by FDA predicate rules or submitted to FDA. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11

FDA Part 11 guidance: FDA guidance: Part 11, Electronic Records; Electronic Signatures - Scope and Application. FDA, September 2003, final guidance; FDA intends to exercise enforcement discretion on the Part 11 validation, audit trail, record retention and record copying requirements; predicate rules still apply. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/part-11-electronic-records-electronic-signatures-scope-and-application

FDA AI draft guidance: FDA draft guidance: Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products. FDA, January 2025, draft, not final. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

OWASP LLM Top 10 2026: OWASP Top 10 for LLM Applications 2026. OWASP GenAI Security Project, 2026 edition with re-ranked IDs. This questionnaire cites the 2025 item pages; each 2025 entry above gives its 2026 number. https://genai.owasp.org/resource/owasp-genai-llm-top-10-2026/

EMA GCP computerised systems guideline: EMA Guideline on computerised systems and electronic data in clinical trials (EMA/INS/GCP/112288/2023). EMA GCP Inspectors Working Group, 2023; the EU counterpart for clinical-trial systems. https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-computerised-systems-electronic-data-clinical-trials_en.pdf

Draft EU GMP Annex 22: EudraLex Volume 4, draft Annex 22: Artificial Intelligence (consultation document). European Commission, consultation draft, not in force; excludes generative AI and large language models, which it says should not be used in critical GMP applications. https://health.ec.europa.eu/document/download/5f38a92d-bb8e-4264-8898-ea076e926db6_en?filename=mp_vol4_chap4_annex22_consultation_guideline_en.pdf

You may copy, edit and use this questionnaire inside your organization and send it to vendors you are assessing.