

A SUPERWISE® LEARN KIT · OCTOBER 2026

Five questions to ask about any AI tool

A worksheet for the person who approves ChatGPT, Claude, Copilot, Gemini, or the AI inside practice software. Ask five questions, keep the evidence, and file the answers with the contract.

Ask them in this order, of every tool

The same five questions work on a chat assistant, an AI feature inside your accounting or practice software, and a tool one team bought on a card. Put them to the vendor in writing, then check the answers against the product itself.

ANSWERED FROM DOCUMENTS The vendor tells you. Ask for the clause.

1 DESTINATION

Where does the data go?

A good answer names where prompts and files are processed and stored, how long they are kept, whether they train a model, which subprocessors receive them, and the contract term behind each. Record every exception. Anthropic keeps safety-flagged content on its commercial plans for up to two years, after the chat itself is deleted [13].

Evasion: "Your data is secure." That answers a question about security. The question was where it goes.

Test: Find the clause behind each answer, and ask who on either side can change it.

2 PROVENANCE

Who built the model?

Every model and AI processor behind the product, including those reached through optional features, connectors, and agents. A single legal product can sit on several providers: Clio's subprocessor list names seven AI processors for Clio Manage [29].

Evasion: "Our own proprietary AI." Most AI products are built on another company's model.

Test: Ask how the vendor learns of a model change, and how and when it tells you.

↓ The first two are claims a vendor makes. The last three **you can check yourself**, with a test account, a made-up client record, and one exported log entry.

ANSWERED FROM THE RUNNING SYSTEM You test it. Keep what the test produced.

3 CONTROL

Is the rule enforced, or written down?

Name the admin setting, contract term, or technical check that holds each rule, and the person who can change it. On ChatGPT Free and Plus, the training opt-out is a setting each user controls, with no organization console above it [1].

Evasion: "All staff are trained on acceptable use." Training is a good thing to have and holds no rule by itself.

Test: Pick one rule. With permission, try to break it using made-up data, and record what happened.

4 VISIBILITY

What can we see while it runs?

The logs and alerts your own people can open: who used which tool, the input, the output, files, and connected services. In Microsoft 365 Copilot, what Purview shows of prompts and responses depends on licensing and configuration [22].

Evasion: "We send a quarterly report." A report describes the past, written by the party being evaluated.

Test: Send a test prompt and check that the event, its content, and any alert appear where you expect.

Can we prove who did what later?

Export one real record and read it: user, timestamp, content, retention, and who can alter it. Coverage varies by product surface. The Claude Enterprise audit export covers 180 days and carries identifiers in place of chat content [17]. ChatGPT Enterprise keeps Compliance Platform logs for 30 days unless you export them [4].

Evasion: "Everything is logged." Logged where, for how long, tied to whom, and editable by whom?

Test: Trace one interaction from last week to a person, and check how long the record will last.

Send the questions in writing

Before we approve [\[tool and plan\]](#) for [\[firm name\]](#), please answer the following in writing, with a link to the clause or document behind each answer.

- 1 Where are our prompts and files processed and stored, how long are they kept, and are they used to train any model? Which subprocessors receive them?
- 2 Which models and AI providers does the product use, including through connectors and agents? How will you tell us when that changes?
- 3 Which of our rules can an administrator enforce in settings, and who can change those settings?
- 4 What can our administrators see of usage, prompts, outputs, and files, without asking you?
- 5 What audit record can we export, what does each entry contain, and how long is it kept?

Reply to [\[name, email\]](#) by [\[date\]](#). File the reply with the worksheet in part 4.

PLAIN ENGLISH · AUDIT LOG

A record, kept by the system, of who did what and when. For an AI tool that means which person or agent sent a request, what it contained, what came back, and what happened next. It is only as good as what it holds and how long it is kept, which is why question 5 asks you to export one entry and read it.

Each question in more depth

- **All five together.** [superwise.ai/learn/ai-governance-vs-theater/](#)
- **One page each.** [/learn/where-does-the-data-go/](#), [/learn/who-built-the-model/](#), [/learn/enforced-or-written-down/](#), [/learn/what-you-can-see-while-it-runs/](#), and [/learn/proving-who-did-what/](#)

Terms by product and plan, as of October 6, 2026

Read from each vendor's public documentation. Where a cell says a commitment was not found, the public pages did not make it; a negotiated contract can. Plans change names and terms: ChatGPT Team became ChatGPT Business on August 29, 2025 [7]. Numbers in brackets point to the sources at the end of this kit.

The data: questions 1 and 2

PLAN	USED FOR TRAINING	RETENTION	WHERE IT IS PROCESSED
OPENAI • CHATGPT			
Free and Plus	Yes, unless the user turns off "Improve the model for everyone." Temporary Chat is excluded [1].	Kept until deleted, then removed within 30 days, with stated exceptions. Temporary Chats up to 30 days [2].	No region commitment found. The privacy policy describes international transfers [9].
Business	No, by default [3].	Kept until users delete chats, then removed within 30 days. No workspace retention setting found [2].	No Business region commitment found [3].
Enterprise	No, by default [3].	Workspace retention controls. Deleted data can remain internally up to 30 days [3, 4].	At-rest residency in listed regions. Inference residency only in the EEA with Switzerland, the US, and the UAE; apps, MCP servers, and web search follow their own terms [5].
ANTHROPIC • CLAUDE			
Free and Pro	When the user allows it, plus flagged content and feedback. Incognito chats are excluded [10].	Deleted chats leave the backend within 30 days. A de-identified training copy can be kept up to five years [11].	No plan residency commitment. The privacy policy permits international transfers [12].
Team	No, by default. Feedback is an exception [14].	Deleted chats leave within 30 days. Flagged content up to two years, classifier scores up to seven, feedback up to five [13].	Stored in the US. Can be routed through the US, Europe, Asia, or Australia unless agreed otherwise [15].
Enterprise	No, by default, with the same exceptions [14].	Indefinite by default. Configurable, with a 30-day minimum; project files are set separately [16].	Stored in the US. No general at-rest residency promise found [15].
MICROSOFT AND GOOGLE			
Microsoft 365 Copilot	Files and communications are not used to train foundation models [19].	Set in Purview. Copies sit in hidden Exchange folders, and holds can delay deletion [20].	Inside the Microsoft 365 service boundary, under the Microsoft 365 commitments [21].
Gemini for Google Workspace	Not used outside your domain without permission [24].	Gemini app history defaults to 18 months; admins choose 3, 18, or 36 months, or off. In-app history behaves separately [25, 26].	Global by default. Data-region controls on specified editions only [27].

The controls: questions 3 to 5

PLAN	ADMIN CONTROLS	AUDIT LOG	BUSINESS ASSOCIATE AGREEMENT
OPENAI · CHATGPT			
Free and Plus	Individual settings only [1].	None for an organization. The Compliance Platform is Enterprise and Edu only [4].	Not eligible [6].
Business	Workspace admin and app controls, with fewer features than Enterprise [8].	Compliance Platform not offered [4].	Not eligible [6].
Enterprise	SAML SSO, feature, app, and retention controls [3].	Append-only Compliance Platform logs, kept 30 days; export to keep longer [4].	In a Regulated Workspace with a BAA, with feature limits [6].
ANTHROPIC · CLAUDE			
Free and Pro	Individual settings only [10].	None for an organization. Audit logs are Enterprise only [17].	Not eligible [18].
Team	Owners manage members, features, and data exports [28].	Not included [17].	Not eligible [18].
Enterprise	SSO, SCIM, feature controls, exports, and custom retention [16].	Owners export 180 days of events. Identifiers only; chat and project content needs a separate export [17].	After HIPAA is enabled and the BAA accepted. Excluded features and third parties stay outside it [18].
MICROSOFT AND GOOGLE			
Microsoft 365 Copilot	Entra permissions, sensitivity labels, DLP, and retention, depending on license [22].	Audit (Standard) records interactions once auditing is on: user, time, and resources referenced [23].	Microsoft's HIPAA BAA, for in-scope services you license and configure [30].
Gemini for Google Workspace	Turn the service on or off, set history, limit data access, and apply DLP where licensed [31].	Eligible editions log Gemini activity. No single content log across every Gemini surface is documented [32].	Under Google's BAA, with some features excluded [24].

Your contract decides. Public documentation cannot show your negotiated terms, your tenant's settings, or which optional features your staff switched on. Use this table to know what to ask for, then file the vendor's own answer with the worksheet.

The five questions, run against real products

A consumer chatbot, an enterprise assistant, and a legal tool, answered from public documentation read October 6, 2026.

CONSUMER CHATBOT

ChatGPT Free

QUESTION	WHAT THE PUBLIC DOCUMENTS SAY	MONDAY ACTION
1 • Data	Chats are kept until deleted, then removed within 30 days, with exceptions. The privacy policy permits international transfers [2, 9].	Keep client and regulated data out until a business plan is approved.
2 • Model	OpenAI supplies the model. GPTs and their actions can send data to third parties under separate terms [33].	List the GPTs, actions, and connected services in use.
3 • Control	The training opt-out and Temporary Chat are choices each user makes. Free has no organization policy controls [1].	Treat a screenshot of one user's setting as weak evidence.
4 • Visibility	Each user sees their own history. No organization monitoring is included [4].	Record it as a visibility gap in the risk register.
5 • Evidence	No organization audit trail ties an interaction to a workforce identity [4].	Keep it out of work that may need legal hold or review.

ENTERPRISE ASSISTANT

Microsoft 365 Copilot

QUESTION	WHAT THE PUBLIC DOCUMENTS SAY	MONDAY ACTION
1 • Data	Prompts, responses, and referenced files stay inside Microsoft 365 services, with retention copies in hidden Exchange folders [20, 22].	Map tenant geography, retention, holds, and agents first.
2 • Model	Microsoft runs Copilot on Azure OpenAI models plus its own orchestration and grounding [21].	Record every plugin, connector, and agent as part of it.
3 • Control	Existing permissions, labels, and DLP apply. Copilot can surface anything the signed-in user can already open [34].	Fix overshared folders before licensing users.
4 • Visibility	Purview can show usage, prompts, responses, and policy events, by license and configuration [35].	Prompt against a labeled file; check the event and alert.
5 • Evidence	Audit (Standard) records user, time, location, and resources referenced [23].	Export a test record; note retention and content fields.

LEGAL PRACTICE SOFTWARE

Clio Manage AI

QUESTION	WHAT THE PUBLIC DOCUMENTS SAY	MONDAY ACTION
1 • Data	Clio says AI data is encrypted and processed in the customer's region. Its subprocessor list scopes each AI processor by product and location [29, 36].	Match your region and features to each subprocessor row.
2 • Model	The list names AWS Bedrock, Anthropic, OpenAI, Google Vertex AI, Azure Foundry, Braintrust, and Turbopuffer [29].	Ask which processors each workflow you use calls.
3 • Control	An admin activates the package. It follows matter permissions, and buying it turns it on for every firm user [37].	Test a restricted matter with two accounts.
4 • Visibility	Source-linked drafts and human review of actions are described. No AI-specific live console was found [37].	Ask for a demonstration; mark it unverified until then.
5 • Evidence	The public pages do not describe an exportable AI audit record. The terms require human oversight [38].	Run one task and export the record it left.

One tool, five answers, one decision

Copy this page once for each tool. Write the answer, note the evidence you hold, and tick how you got it. An answer you were only told is a claim; an answer you tested is evidence.

TOOL AND PLAN	ACCOUNT OR WORKSPACE	USED BY
REVIEWED BY	DATE	CONTRACT OR ORDER FORM

1 Where does the data go? Processing and storage location, retention, training use, subprocessors, exceptions.

ANSWER

CLAUSE OR DOCUMENT

☐ Shown in writing ☐ Told only ☐ No answer

2 Who built the model? Every model and AI processor, including those behind connectors and agents.

ANSWER

CLAUSE OR DOCUMENT

☐ Shown in writing ☐ Told only ☐ No answer

3 Is the rule enforced, or written down? The setting, term, or check that holds each rule, and who can change it.

ANSWER

WHAT THE TEST PRODUCED

☐ Tested and held ☐ Told only ☐ Failed or no answer

4 What can we see while it runs? User, input, output, files, and connected services, on a screen your own people open.

ANSWER

WHAT THE TEST PRODUCED

☐ Tested and held ☐ Told only ☐ Failed or no answer

5 Can we prove who did what later? One exported record: user, timestamp, content, retention, and who can alter it.

ANSWER

WHERE THE EXPORT IS FILED

☐ Tested and held ☐ Told only ☐ Failed or no answer

Decision

☐ Approve ☐ Approve with limits ☐ Reject

Data allowed in this tool _____

Limits or settings required _____

Approved by _____ Next review _____

What to file with each approval

The worksheet is the summary. These are the documents behind it, so that whoever reviews the tool next year starts from evidence.

Contract

- ☐ **The plan name and order form.** Plan names change. Record the name on the invoice and the date you checked it.
- ☐ **The data processing agreement, and the BAA where health information is involved.** Note who accepted each and the account it covers.
- ☐ **The subprocessor list, as of the approval date.** Save a copy; vendors update the page in place.

Configuration

- ☐ **Screenshots of the admin settings.** Training use, retention, connectors, and any feature you switched off.
- ☐ **The residency selection, if the plan offers one.** Note the features the vendor excludes from it.
- ☐ **The features the contract excludes.** List each one with the setting that keeps it off.

Evidence

- ☐ **One exported audit event.** The record from question 5, with the date you exported it.
- ☐ **The result of a deletion test.** Delete a test chat and record what the vendor's tools show afterwards.
- ☐ **The source pages and the date you read them.** The URLs behind each answer, so a reviewer can re-read them.

Recheck the five questions when

- **The plan changes.** A rename, an upgrade, or a downgrade moves the terms with it.
- **The contract renews, or the vendor sends a notice of change.** Read what changed against the worksheet.
- **Someone connects something new.** A connector, an agent, or a browser extension reaches data the original approval did not cover.

PLAIN ENGLISH • SUBPROCESSOR

Another company your vendor hands your data to so it can deliver the service, such as the AI model provider behind a legal or accounting tool. Your data goes wherever the subprocessors take it, so their list is part of the answer to question 1.

Answers you hold yourself, for the traffic you route

Questions 3 to 5 are the hard ones on a consumer plan, because the vendor gives the firm no console. A gateway on the PC moves part of each answer onto equipment the firm controls.

What SUPERWISE Sentinel does

As described on superwise.ai/sentinel/, read October 8, 2026 [39]:

- **Placeholders before the prompt leaves.** SUPERWISE® Sentinel replaces emails, phone numbers, account and card numbers, government IDs, and API keys with placeholders before a prompt reaches ChatGPT or Claude. That is a tested answer to question 3 for those data types.
- **A record you can open.** Every inspected interaction lands in an Activity Log: what was checked, what was redacted, and when. That is evidence for questions 4 and 5.
- **No new destination.** Sentinel runs on the machine or in your own cloud, and nothing is routed through SUPERWISE. The answer to question 1 stays the vendor's answer.

Limits to write on the worksheet

- **It covers the traffic sent through it.** On Windows that is Claude Code, Claude Desktop, and, through a browser extension in preview, ChatGPT and Claude in Chrome and Edge. A tool or device outside that path keeps the answers in parts 2 and 3.
- **It catches the data types it lists.** Names and street addresses are absent from that list, so a prompt can still carry a client's identity.
- **It leaves the vendor's terms as they are.** Retention, training use, and any BAA still come from the vendor's contract. Sentinel makes no tool compliant.

Trying it. SUPERWISE Sentinel installs on Windows 10 and 11 without admin rights. It is free for 30 days with no card, then \$10, \$25, or \$50 a month per deployment; the \$50 Business plan is the one listed for staff using ChatGPT and Claude with client data, as shown on superwise.ai/pricing/ on October 8, 2026 [39].

Sources

Vendor documentation read October 6, 2026, unless marked. Anthropic, Microsoft, and Google pages were checked again on October 8, 2026. Vendors edit these pages in place, so re-read a source before you rely on it.

- 1 OpenAI Help Center, *Data Controls FAQ*. help.openai.com/en/articles/7730893
- 2 OpenAI Help Center, *How do I opt out of my data being used to train future models?* Retention and deletion. help.openai.com/en/articles/8983778
- 3 OpenAI, *Enterprise privacy*. Business and Enterprise training, retention, and admin commitments. openai.com/enterprise-privacy/
- 4 OpenAI Help Center, *Compliance API for Enterprise customers*. Plans, log retention, append-only logs. help.openai.com/en/articles/9261474
- 5 OpenAI Help Center, *Data residency and inference residency for ChatGPT*. help.openai.com/en/articles/9903489
- 6 OpenAI Help Center, *HIPAA eligible products and functionality*. help.openai.com/en/articles/20001069
- 7 OpenAI Help Center, *ChatGPT Team is now ChatGPT Business*. help.openai.com/en/articles/12111915
- 8 OpenAI Help Center, *Admin controls, security, and compliance in apps*. help.openai.com/en/articles/11509118
- 9 OpenAI, *Privacy policy*. International transfers. openai.com/policies/row-privacy-policy/
- 10 Anthropic Privacy Center, *Is my data used for model training?* (consumer plans). privacy.claude.com/en/articles/10023580
- 11 Anthropic Privacy Center, *How long do you store my data?* (consumer plans). privacy.claude.com/en/articles/10023548
- 12 Anthropic, *Privacy policy*. International transfers. anthropic.com/legal/privacy
- 13 Anthropic Privacy Center, *How long do you store my organization's data?* Checked October 8, 2026. privacy.claude.com/en/articles/7996866
- 14 Anthropic Privacy Center, *Is my data used for model training?* (commercial plans). privacy.claude.com/en/articles/7996868
- 15 Anthropic Privacy Center, *Where are your servers located?* privacy.claude.com/en/articles/7996890
- 16 Anthropic Privacy Center, *Configure custom data retention controls for Enterprise plans*. Checked October 8, 2026. privacy.claude.com/en/articles/10440198
- 17 Claude Help Center, *Access audit logs*. 180-day export, identifiers only. Checked October 8, 2026. support.claude.com/en/articles/9970975
- 18 Anthropic Privacy Center, *Business Associate Agreements (BAA) for Commercial Customers*. privacy.claude.com/en/articles/8114513
- 19 Microsoft Learn, *Microsoft 365 Copilot minimum requirements and data compliance*. learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-minimum-requirements-data-compliance
- 20 Microsoft Learn, *Retention policies for Copilot*. learn.microsoft.com/en-us/purview/retention-policies-copilot
- 21 Microsoft Learn, *Microsoft 365 Copilot architecture*. learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-architecture
- 22 Microsoft Learn, *Data protection and auditing architecture for Microsoft 365 Copilot*. learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-architecture-data-protection-auditing
- 23 Microsoft Learn, *Audit logs for Copilot and AI applications*. Checked October 8, 2026. learn.microsoft.com/en-us/purview/audit-copilot
- 24 Google Workspace Admin Help, *Gemini for Google Workspace FAQ*. Training use and BAA. support.google.com/a/answer/14130944
- 25 Google Workspace Admin Help, *Gemini app settings for admins*. History periods. Checked October 8, 2026. support.google.com/a/answer/14571493
- 26 Google Help, *Gemini conversation history in Workspace apps*. support.google.com/meet/answer/16880047
- 27 Google Workspace Admin Help, *Choose a geographic location for your data*, and the data-region editions list. support.google.com/a/answer/14310028 · /14316863
- 28 Anthropic Privacy Center, *Who owns and manages the data of my Team?* privacy.claude.com/en/articles/9265372
- 29 Clio, *Subprocessors*. AI processors for Clio Manage. clio.com/tos/subprocessors/
- 30 Microsoft Learn, *HIPAA and HITECH offering*. learn.microsoft.com/en-us/compliance/regulatory/offering-hipaa-hitech
- 31 Google Workspace Help, *Gemini access controls*. support.google.com/a/users/answer/17010577
- 32 Google Workspace Admin Help, *Generative AI Privacy Hub*. support.google.com/a/answer/15706919
- 33 OpenAI Help Center, *GPTs data privacy FAQs*. help.openai.com/en/articles/8554402
- 34 Microsoft Learn, *Security for Microsoft 365 Copilot*. learn.microsoft.com/en-us/microsoft-365/copilot/security-microsoft-365-copilot
- 35 Microsoft Learn, *Copilot Control System: security and governance*. learn.microsoft.com/en-us/microsoft-365/copilot/copilot-control-system/security-governance
- 36 Clio, *Security*. clio.com/security/
- 37 Clio Help Center, *What is Manage AI and how to access it*. help.clio.com/hc/en-us/articles/29657814904475
- 38 Clio, *Terms of service*, section 14. clio.com/tos/
- 39 SUPERWISE, *Sentinel and Pricing*, read October 8, 2026. Every product statement in part 6. superwise.ai/sentinel/ · superwise.ai/pricing/