

AI Vendor Packet sample report

A sample customer commitment drift packet for security, privacy, legal, and founder review. Internal review only, not legal advice.

GENERATED 2026-05-24	STATUS Needs owner review
VENDORS 5	COMMITMENTS 6
TOP FINDINGS 8	FILES PDF and CSV

Reader path

- 1. Start with the conclusion**
Read the executive summary, scope conclusion, and review status before opening the appendices.
- 2. Route the open decisions**
Use the first decisions and action plan to assign the owners who must approve, qualify, or hold customer answers.
- 3. Use sources as evidence**
Open the source citations, source excerpts, and evidence table when a reviewer asks what supports a finding.
- 4. Keep the boundary attached**
Attach limitations, the not-legal-advice note, and the readiness checks when the PDF or CSV is routed internally.

Boundary
AI Vendor Packet organizes official-source review evidence and suggested review actions. It does not provide legal advice, vendor certification, or final approval for customer answers.

Report cover sheet

Public sample report: Sample AI vendor commitment review report

Report ID: sample-report

Status: Needs owner review

Audience: Security / GRC, Legal / Privacy, and Founder / Operator reviewers

Primary use: Route vendor commitment review work with official source links before Trust Center, DPA, questionnaire, SOC 2, or customer response language goes out.

External use status: Sample only. Use it to inspect the report structure, not as current vendor coverage for a workspace.

Scope snapshot

- Report type: Public sample report
- Workspace: Sample workspace
- Period: Illustrative sample
- Generated: 2026-05-24
- Source coverage: Reviewable with visible gaps
- Scope: 5 vendors, 6 commitments, 8 source documents, and 8 material findings are represented.
- Open work: 8 open findings and 0 source coverage gaps should be reviewed before external use.

Boundary

AI Vendor Packet organizes official-source review evidence and suggested review actions. It does not provide legal advice, vendor certification, or final approval for customer answers.

Executive summary

Needs owner review: Use this sample to inspect the report structure, not as current coverage for your company.

5 vendors and 6 commitments are shown with 8 review prompts so your team can judge whether the paid report is useful. The findings are illustrative and should not be treated as live vendor coverage.

What changed

- The sample highlights AI data-use, retention, DPA, subprocessor, and security statements that often affect customer commitments.

-
- 8 sample findings show where broad vendor language may need product-path review before external use.
 - 23 source citations show the official-source trail a paid report should preserve.

Why it matters

- Trust Center answers, DPA exhibits, security questionnaire responses, and SOC 2 monitoring evidence packet notes should stay aligned with the exact vendor source and product path.
- 7 high-priority items should be reviewed before relying on broad AI data-use or monitoring evidence language externally.
- 8 decision records identify who should approve, edit, hold, or mark the language not applicable.

Next review actions

- Resolve failed, stale, or not-checked sources before using the report as evidence for customer answers.
- Confirm the exact product path, plan, data category, region, and customer commitment language for each material finding.
- Attach the source citations and decision register to the internal review ticket or SOC 2 monitoring evidence packet.

Do not conclude

- Do not treat this report as legal advice, vendor certification, or final customer-facing approval.
- Do not treat sample findings as current paid coverage for your workspace.
- Do not assume unchecked, stale, or failed sources are safe to cite externally.

Evidence snapshot

- 8 source documents checked or represented in the report scope.
- 8 decision records and 23 source citations are available for review.
- 0 captured excerpts are present; findings without captured excerpts keep the official URL and evidence limit visible.

Conclusion for this scope

Internal review only: For a production workflow using customer content, personal data, EU data, this sample shows how an internal AI-vendor commitment review should be organized. Do not copy customer wording until the named owners have made the scope

decisions.

Scope basis

- Selected vendors: OpenAI, Anthropic, Azure OpenAI / Microsoft AI, Google Vertex AI / Gemini for Cloud, AWS Bedrock.
- Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.
- Commitment scope: 6 selected commitments plus the custom AI-training commitment.
- Evidence basis: 23 source citations, 8 findings, and 8 source notes or excerpts.

Immediate actions

- Resolve the first 3 decisions before using no-training, retention, DPA, or Trust Center language externally.
- Attach the source citations and decision register to the internal review ticket or SOC 2 monitoring evidence packet.
- Treat unresolved product-path, agreement, model, region, and downstream storage questions as holds, not as clean evidence.

Limit: Do not use this sample conclusion as legal advice, vendor certification, or final approval for a Trust Center answer, DPA exhibit, security questionnaire, or customer response.

Decisions to resolve first

- 1. OpenAI: Does the customer answer describe OpenAI API Platform use only, or does it also cover ChatGPT workspace or unmanaged account use?

Owner: Security / GRC; evidence: Data controls in the OpenAI platform; reviewed May 21, 2026.

Needed before external use: Answer the applicability question before using this finding

externally: Does the customer answer describe OpenAI API Platform use only, or does it also cover ChatGPT workspace or unmanaged account use?

External use limit: Hold external answers until product path, agreement scope, and applicability are resolved.

- 2. Anthropic: Is a zero data retention agreement actually signed, and which Anthropic products does it cover?

Owner: Legal / Privacy; evidence: Zero data retention agreement applicability; reviewed May 21, 2026.

Needed before external use: Answer the applicability question before using this finding

externally: Is a zero data retention agreement actually signed, and which Anthropic products does it cover?

External use limit: Hold external answers until product path, agreement scope, and applicability are resolved.

- 3. Azure OpenAI / Microsoft AI: Should the customer DPA exhibit name Microsoft, OpenAI, or both for this workflow?

Owner: Legal / Privacy; evidence: Microsoft Products and Services Data Protection Addendum; reviewed May 21, 2026.

Needed before external use: Answer the applicability question before using this finding externally: Should the customer DPA exhibit name Microsoft, OpenAI, or both for this workflow?

External use limit: Hold external answers until product path, agreement scope, and applicability are resolved.

Action plan

Sample action plan: Use this plan to inspect how a paid report should move from source evidence to owner decisions and closeout.

Turn the report into a recorded internal review decision before Trust Center, DPA, security questionnaire, audit, or customer-response language is used externally.

Closure criteria

- Scope confirmed: Report cover sheet, selected vendors, customer commitments, and period. Pass: A reviewer can tell which vendors, product paths, data categories, and customer-facing commitments are in scope.
- Source evidence current enough: Source citations, source coverage, and reviewed dates. Pass: Failed, stale, and not-checked sources are either refreshed, accepted as explicit limitations, or held out of external use.
- Owner decisions recorded: Decision register and reviewer notes. Pass: Each material finding has an owner decision, applicability decision, or documented hold.
- Files archived: PDF, CSV, and 23 source citations. Pass: The report files are attached to the internal review ticket, SOC 2 record, or legal/privacy review file.
- External-use decision recorded: External use status and closeout note. Pass: A named owner records whether the report can support a Trust Center, DPA, questionnaire, audit, or customer-response update.

Review sequence

- 1. Founder / Operator: Confirm report scope and external use limit; Sample only. Evidence: Cover sheet, Reviewable with visible gaps, and Needs owner review. Completion: Scope is confirmed or corrected before the report is routed as review evidence.

External use check: Sample only. Use it to inspect the report structure, not as current vendor coverage for a workspace.

- 2. Security / GRC: Resolve source freshness and coverage gaps; Sample only.

Evidence: 0 failed source checks, 0 stale source(s), 15 not-checked registry sources. 23 source citations are available.

Completion: Source gaps are refreshed, copied into limitations, or marked as a hold before external use.

External use check: Hold external answers for affected findings until source gaps are resolved.

- 3. Security / GRC: Answer the applicability question; Sample only.

Evidence: OpenAI: Data controls in the OpenAI platform; reviewed May 21, 2026; Fresh.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally: Does the customer answer describe OpenAI API Platform use only, or does it also cover ChatGPT workspace or unmanaged account use?

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 4. Security / GRC: Answer the applicability question; Sample only.

Evidence: OpenAI: OpenAI Data Processing Addendum; reviewed May 21, 2026; Fresh.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally:

Confirm exact product path, plan, data categories, region, and customer commitment language before using this as evidence for customer answers.

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 5. Security / GRC: Answer the applicability question; Sample only.

Evidence: Anthropic: Is my data used for model training?; reviewed May 21, 2026; Fresh.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally:

Confirm exact product path, plan, data categories, region, and customer commitment language before using this as evidence for customer answers.

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 6. Legal / Privacy: Answer the applicability question; Sample only.

Evidence: Anthropic: Zero data retention agreement applicability; reviewed May 21, 2026;

Fresh. Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally: Is a zero data retention agreement actually signed, and which Anthropic products does it cover?

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 7. Security / GRC: Answer the applicability question; Sample only.

Evidence: Azure OpenAI / Microsoft AI: Data, privacy, and security for Models sold by Azure in Microsoft Foundry; reviewed May 21, 2026; Fresh. Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally:

Confirm exact product path, plan, data categories, region, and customer commitment language before using this as evidence for customer answers.

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 8. Security / GRC: Answer the applicability question; Sample only.

Evidence: Google Vertex AI / Gemini for Cloud: Vertex AI zero data retention; reviewed May 21, 2026; Fresh. Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally: Is the selected Vertex AI model, endpoint, and feature eligible for the retention language being reused?

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 9. Founder / Operator: Answer the applicability question; Sample only.

Evidence: AWS Bedrock: Data protection in Amazon Bedrock; reviewed May 21, 2026; Fresh.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally:

Which Bedrock models, regions, logging paths, agents, and knowledge bases are enabled?

External use check: Hold external answers until product path, agreement scope, and applicability are resolved.

- 10. Legal / Privacy: Answer the applicability question; Sample only.

Evidence: Azure OpenAI / Microsoft AI: Microsoft Products and Services Data Protection Addendum; reviewed May 21, 2026; Fresh. Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Completion: Answer the applicability question before using this finding externally:

Should the customer DPA exhibit name Microsoft, OpenAI, or both for this workflow?

External use check: Hold external answers until product path, agreement scope, and

applicability are resolved.

Escalation triggers

- An official vendor source is failed, stale, missing, or outside the reviewed product path.
- A finding depends on plan, region, agreement, model, or data category scope that is not recorded.
- A customer answer would use vendor language before the named owner records a decision.
- A reviewer wants to treat the report as legal advice, vendor certification, or final external approval.
- Record the exact vendor products, plans, and data categories before reusing these answers with customers.
- Attach official source links and a last reviewed date to each Trust Center or questionnaire answer.
- Review DPA and subprocessor evidence before answering personal data questions.

Closeout note

Close the report only after source gaps, applicability questions, owner decisions, file archival, and external use limits are recorded.

Review readiness checks

Sample only: Sample findings are illustrative and cannot be accepted as current workspace coverage.

This sample shows the report format only; it is not evidence for customer answers.

A proceed status means the report is organized for internal review. It does not provide legal advice, vendor certification, or final approval for customer answers.

Readiness checks

- Scope and identity: Sample only.

Review question: Can a reviewer see the report scope before reading findings?

Evidence: Public sample report; Report type: Public sample report Workspace: Sample workspace Period: Illustrative sample Generated: 2026-05-24 Source coverage: Reviewable with visible gaps Scope: 5 vendors, 6 commitments, 8 source documents, and 8 material findings are represented. Open work: 8 open findings and 0 source coverage gaps should be reviewed before external use.

Rule: Report type, workspace, period, source coverage, vendors, commitments, and open work are named.

Next step: Keep the evidence attached to the report record.

- Source coverage ready: Sample only.

Review question: Are failed, stale, and not-checked official sources resolved or held?

Evidence: 8 checked; 0 failed; 0 stale; 15 not checked.

Rule: No failed, stale, or not-checked source is treated as ready customer evidence.

Next step: Refresh the source, mark it out of scope, or keep the affected row on hold.

- Evidence traceability: Sample only.

Review question: Can each material row be traced to source citations and evidence status?

Evidence: 23 source citations; 15 re-checks.

Rule: Finding evidence and coverage-only sources have official URLs, reviewed dates, evidence status, and boundaries.

Next step: Attach citations and resolve any re-checks before closeout.

- Owner decision record: Sample only.

Review question: Has each material finding been routed to an owner decision?

Evidence: 8 decision records; 8 open findings; 8 unresolved decisions.

Rule: Each material finding has a reviewer note, applicability decision, or documented hold.

Next step: Record approve, edit, not-applicable, or hold for every material finding.

- Applicability assumptions: Sample only.

Review question: Are product path, plan, data categories, and agreement assumptions clear?

Evidence: 0 recorded; 8 unresolved; 0 source re-check.

Rule: Applicability is recorded or explicitly held before external use.

Next step: Confirm product path, plan, region, agreement scope, data categories, and customer commitment wording.

- Action plan: Sample only.

Review question: Does the packet give reviewers a closure path?

Evidence: Sample action plan; 12 action steps; 5 closure criteria.

Rule: Owners, evidence to use, completion criteria, escalation triggers, and closeout note are present.

Next step: Complete action-plan holds and record the closeout decision.

- Files and archive path: Sample only.

Review question: Are the PDF and CSV ready to attach to the system of record?

Evidence: PDF available: Use the PDF as the review packet for security, legal, privacy, compliance, or founder review. CSV available: Use the CSV to assign findings, track review status, preserve source links, and copy evidence into internal systems.

Rule: PDF and CSV are available, and system-of-record targets are named for archival.

Next step: Keep the evidence attached to the report record.

Applicability checks

This sample shows which product path, data category, agreement, and commitment assumptions a paid report should make explicit.

Applicability counts: 0 recorded; 8 unresolved; 0 source re-check.

Closure rule: Close applicability only when product path, plan, model, region, data categories, agreement scope, and customer commitment wording are recorded or explicitly held.

Applicability details

- OpenAI: Confirm OpenAI API data-use scope before reusing customer AI training language; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- OpenAI: Separate provider retention from your own application retention; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- Anthropic: Tie Claude model-training answers to the product path actually used; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- Anthropic: Verify whether zero data retention is actually covered by your Anthropic agreement; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement,

and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- Azure OpenAI / Microsoft AI: Keep Azure OpenAI evidence separate from direct OpenAI evidence; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- Google Vertex AI / Gemini for Cloud: Confirm Vertex AI zero data retention before using that phrase with customers; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- AWS Bedrock: Use Bedrock-specific evidence for model provider access questions; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

- Azure OpenAI / Microsoft AI: Review whether customer DPA exhibits should name Microsoft, OpenAI, or both; Sample only.

Product path and plan: Customer-facing AI features using direct API and cloud-hosted model providers.

Data categories: customer content, personal data, EU data; usage stage: production

Required confirmation: Confirm the real workspace product path, plan, region, agreement, and data categories before using this sample as evidence.

External use limit: Do not use sample assumptions as evidence in customer answers.

Applicability limit

Applicability checks are review prompts and owner-confirmation records. They are not legal advice, vendor certification, or automatic approval for customer answers.

Decision register

Decision counts: 0 recorded; 0 owner decision; 8 applicability decision; 0 source re-check.

- OpenAI: Confirm OpenAI API data-use scope before reusing customer AI training language; Needs applicability decision.

Owner: Security / GRC; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Does the customer answer describe OpenAI API Platform use only, or does it also cover ChatGPT workspace or unmanaged account use?

- OpenAI: Separate provider retention from your own application retention; Needs applicability decision.

Owner: Security / GRC; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Confirm exact product path, plan, data categories, region, and customer commitment language before using this as evidence for customer answers.

- Anthropic: Tie Claude model-training answers to the product path actually used; Needs applicability decision.

Owner: Security / GRC; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Confirm exact product path, plan, data categories, region, and customer commitment language before using this as evidence for customer answers.

- Anthropic: Verify whether zero data retention is actually covered by your Anthropic agreement; Needs applicability decision.

Owner: Legal / Privacy; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Is a zero data retention agreement actually signed, and which Anthropic products does it cover?

- Azure OpenAI / Microsoft AI: Keep Azure OpenAI evidence separate from direct OpenAI evidence; Needs applicability decision.

Owner: Security / GRC; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Confirm exact product path, plan, data categories, region, and customer commitment language before using this as evidence for customer answers.

- Google Vertex AI / Gemini for Cloud: Confirm Vertex AI zero data retention before using that phrase with customers; Needs applicability decision.

Owner: Security / GRC; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Is the selected Vertex AI model, endpoint, and feature eligible for the retention language being reused?

- AWS Bedrock: Use Bedrock-specific evidence for model provider access questions; Needs applicability decision.

Owner: Founder / Operator; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Which Bedrock models, regions, logging paths, agents, and knowledge bases are enabled?

- Azure OpenAI / Microsoft AI: Review whether customer DPA exhibits should name Microsoft, OpenAI, or both; Needs applicability decision.

Owner: Legal / Privacy; external use: Hold external answers until product path, agreement scope, and applicability are resolved.

Record action: Answer the applicability question before using this finding externally:

Should the customer DPA exhibit name Microsoft, OpenAI, or both for this workflow?

Coverage and gaps

Sample shape, not paid coverage: This public sample shows the report structure using selected vendors and sources. It is not current paid coverage for your workspace.

Completeness label: Reviewable with visible gaps. 8 sources are checked against 23 expected sources; 15 not checked source issues remain visible before external use.

Coverage counts: 8 checked; 0 failed; 0 stale; 15 not checked.

Source coverage by vendor

- OpenAI: Reviewable with visible gaps; 2 checked, 0 failed, 0 stale, 3 not checked.
- Anthropic: Reviewable with visible gaps; 2 checked, 0 failed, 0 stale, 4 not checked.
- Azure OpenAI / Microsoft AI: Reviewable with visible gaps; 2 checked, 0 failed, 0 stale, 2 not checked.
- Google Vertex AI / Gemini for Cloud: Reviewable with visible gaps; 1 checked, 0 failed, 0 stale, 4 not checked.
- AWS Bedrock: Reviewable with visible gaps; 1 checked, 0 failed, 0 stale, 2 not checked.

Checked and not checked sources

- Anthropic: Commercial Terms of Service; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- Anthropic: Custom Data Retention Controls for Claude Enterprise; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- Anthropic: Data Processing Addendum; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- Anthropic: Privacy Policy; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- Anthropic: Is my data used for model training?; Checked; Checked in this packet. Reviewed 1 day ago. This is recent enough for the current packet freshness model.
- Anthropic: Zero data retention agreement applicability; Checked; Checked in this packet. Reviewed 1 day ago. This is recent enough for the current packet freshness model.
- AWS Bedrock: AWS Service Terms; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- AWS Bedrock: Cloud Security; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- AWS Bedrock: Data protection in Amazon Bedrock; Checked; Checked in this packet. Reviewed 1 day ago. This is recent enough for the current packet freshness model.
- Azure OpenAI / Microsoft AI: Microsoft Azure Product Terms; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- Azure OpenAI / Microsoft AI: Microsoft Privacy Statement; Not checked; Verified official source exists in the registry, but this public sample only shows selected report items. A paid scoped report should include it, explicitly exclude it, or show the reason it was not checked.
- Azure OpenAI / Microsoft AI: Data, privacy, and security for Models sold by Azure in Microsoft Foundry; Checked; Checked in this packet. Reviewed 1 day ago. This is recent enough for the current packet freshness model.

Source citations

Citation counts: 8 finding evidence; 15 coverage source; 0 captured excerpt; 15 re-check needed.

- OpenAI: Data controls in the OpenAI platform; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Confirm OpenAI API data-use scope before reusing customer AI training language; evidence status: Source note

URL: <https://platform.openai.com/docs/guides/your-data>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- OpenAI: OpenAI Data Processing Addendum; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Separate provider retention from your own application retention; evidence status: Sample link only

URL: <https://openai.com/policies/data-processing-addendum/>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Anthropic: Is my data used for model training?; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Tie Claude model-training answers to the product path actually used; evidence status: Sample link only

URL:

<https://privacy.anthropic.com/en/articles/7996868-i-want-to-opt-out-of-my-prompts-and-results-being-used-for-training-models>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Anthropic: Zero data retention agreement applicability; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Verify whether zero data retention is actually covered by your Anthropic agreement; evidence status: Source note

URL:

<https://privacy.anthropic.com/en/articles/8956058-i-have-a-zero-data-retention-agreement-with-anthropic-what-products-does-it-apply-to>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Azure OpenAI / Microsoft AI: Data, privacy, and security for Models sold by Azure in Microsoft Foundry; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Keep Azure OpenAI evidence separate from direct OpenAI evidence; evidence status: Sample link only

URL:

<https://learn.microsoft.com/en-us/azure/ai-foundry/responsible-ai/openai/data-privacy>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Google Vertex AI / Gemini for Cloud: Vertex AI zero data retention; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Confirm Vertex AI zero data retention before using that phrase with customers; evidence status: Source note

URL:

<https://docs.cloud.google.com/vertex-ai/generative-ai/docs/vertex-ai-zero-data-retention>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- AWS Bedrock: Data protection in Amazon Bedrock; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Use Bedrock-specific evidence for model provider access questions; evidence status: Source note

URL: <https://docs.aws.amazon.com/bedrock/latest/userguide/data-protection.html>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Azure OpenAI / Microsoft AI: Microsoft Products and Services Data Protection Addendum; Finding evidence; Checked; reviewed May 21, 2026.

Finding: Review whether customer DPA exhibits should name Microsoft, OpenAI, or both; evidence status: Source note

URL:

<https://www.microsoft.com/licensing/docs/view/Microsoft-Products-and-Services-Data-Protection-Addendum-DPA>

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Anthropic: Commercial Terms of Service; Coverage source; Not checked; reviewed May 21, 2026.

Finding: No material finding tied to this source; evidence status: Not checked in report

URL: <https://www.anthropic.com/legal/commercial-terms>

Evidence limit: Coverage evidence only: this source shows traceability, not a finding conclusion.

- Anthropic: Custom Data Retention Controls for Claude Enterprise; Coverage source; Not checked; reviewed May 21, 2026.

Finding: No material finding tied to this source; evidence status: Not checked in report

URL:

<https://privacy.anthropic.com/en/articles/10440198-custom-data-retention-controls-for-claude-enterprise>

Evidence limit: Coverage evidence only: this source shows traceability, not a finding conclusion.

- Anthropic: Data Processing Addendum; Coverage source; Not checked; reviewed May 21, 2026.

Finding: No material finding tied to this source; evidence status: Not checked in report

URL: <https://www.anthropic.com/legal/data-processing-addendum>

Evidence limit: Coverage evidence only: this source shows traceability, not a finding conclusion.

- Anthropic: Privacy Policy; Coverage source; Not checked; reviewed May 21, 2026.

Finding: No material finding tied to this source; evidence status: Not checked in report

URL: <https://www.anthropic.com/legal/privacy>

Evidence limit: Coverage evidence only: this source shows traceability, not a finding conclusion.

Source excerpts and notes

- OpenAI: Confirm OpenAI API data-use scope before reusing customer AI training language; Source note; reviewed May 21, 2026.

Source note or excerpt: OpenAI API Platform data-control evidence is the source to verify model-training and retention scope for API usage.

How to use this source: Check the API Platform path, organization settings, and any separate ChatGPT workspace use before reusing no-training language.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- OpenAI: Separate provider retention from your own application retention; Official source linked; reviewed May 21, 2026.

Source note or excerpt: OpenAI Data Processing Addendum is linked as the official source for this sample finding.

How to use this source: Verify the linked official source before relying on this finding outside the sample report.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- Anthropic: Tie Claude model-training answers to the product path actually used; Official source linked; reviewed May 21, 2026.

Source note or excerpt: Is my data used for model training? is linked as the official source for this sample finding.

How to use this source: Verify the linked official source before relying on this finding outside the sample report.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- Anthropic: Verify whether zero data retention is actually covered by your Anthropic agreement; Source note; reviewed May 21, 2026.

Source note or excerpt: Anthropic retention evidence needs agreement and product-scope validation before a zero-retention claim is reused.

How to use this source: Confirm whether the zero data retention agreement is signed and which product path it covers.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- Azure OpenAI / Microsoft AI: Keep Azure OpenAI evidence separate from direct OpenAI evidence; Official source linked; reviewed May 21, 2026.

Source note or excerpt: Data, privacy, and security for Models sold by Azure in Microsoft Foundry is linked as the official source for this sample finding.

How to use this source: Verify the linked official source before relying on this finding outside the sample report.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- Google Vertex AI / Gemini for Cloud: Confirm Vertex AI zero data retention before using that phrase with customers; Source note; reviewed May 21, 2026.

Source note or excerpt: Vertex AI retention evidence should be checked against the selected model, endpoint, logging path, and downstream storage.

How to use this source: Verify the specific Vertex AI feature and storage path before reusing retention or Trust Center language.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- AWS Bedrock: Use Bedrock-specific evidence for model provider access questions; Source note; reviewed May 21, 2026.

Source note or excerpt: Bedrock data-protection evidence should be reviewed by model, region, logging, agents, and knowledge-base configuration.

How to use this source: Separate Bedrock platform evidence from direct model-provider terms before updating subprocessor or AI-vendor review records.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

- Azure OpenAI / Microsoft AI: Review whether customer DPA exhibits should name Microsoft, OpenAI, or both; Source note; reviewed May 21, 2026.

Source note or excerpt: Microsoft privacy and DPA sources are the controlling evidence for Azure-hosted model workflows and EU-data transfer review.

How to use this source: Confirm whether the customer exhibit should name Microsoft, OpenAI, or both for the actual deployment path.

Evidence limit: Source note only: this is not a legal quote, vendor certification, or final approval for customer answers.

Source freshness

For packet evidence, critical AI and SaaS vendor sources should show a recent reviewed date. Material vendor notices, Trust Center updates, DPA changes, subprocessor notices, and customer-reported changes should be checked before the packet is reused externally.

- Official-source boundary: Source evidence should come from official vendor documentation, Trust Center pages, product documentation, DPAs, or clearly identified primary vendor notices.
- Recent review date: Sources used in a paid packet should have a visible reviewed date and should be rechecked before they are reused for a new customer answer.
- Urgent-change handling: Material vendor notices, broken source links, DPA updates, subprocessor notices, and customer-reported source changes should be routed to the relevant owner before reuse.
- Stale-source warning: A source older than 60 days, missing a reviewed date, or failing the latest source check should be marked for review before the packet is reused externally.
- OpenAI: Data controls in the OpenAI platform; reviewed May 21, 2026; Fresh.
- OpenAI: OpenAI Data Processing Addendum; reviewed May 21, 2026; Fresh.
- Anthropic: Is my data used for model training?; reviewed May 21, 2026; Fresh.
- Anthropic: Zero data retention agreement applicability; reviewed May 21, 2026; Fresh.
- Azure OpenAI / Microsoft AI: Data, privacy, and security for Models sold by Azure in Microsoft Foundry; reviewed May 21, 2026; Fresh.

Selected vendors

OpenAI, Anthropic, Azure OpenAI / Microsoft AI, Google Vertex AI / Gemini for Cloud, AWS Bedrock

Commitment profile

- Customer data is not used for model training.
- Subprocessors are reviewed before material changes.
- Critical vendors have a current review packet.
- Customer data is retained only as long as necessary.
- Trust Center and questionnaire answers must remain accurate.
- EU personal data must be handled under approved transfer mechanisms.
- We do not use customer content to train third-party AI provider models unless the customer has explicitly approved that use.

Claim and evidence table

- OpenAI: Confirm OpenAI API data-use scope before reusing customer AI training language

Commitment: Customer data is not used for model training.

Evidence used: OpenAI: Data controls in the OpenAI platform; reviewed May 21, 2026;

Fresh.

Source evidence recorded: Source note: OpenAI API Platform data-control evidence is the source to verify model-training and retention scope for API usage.

Why this may matter: OpenAI API commitments should cite the platform data controls source and should not be copied to ChatGPT workspace or unmanaged account use without a separate review.

Unknowns: Whether your API organization uses modified retention or abuse monitoring settings.

Owner decision: Security / GRC: Confirm product path, organization settings, and whether prompts are copied elsewhere. Current handoff status: Needs product-path confirmation.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- OpenAI: Separate provider retention from your own application retention

Commitment: No specific customer commitment is mapped yet; treat this as a general review prompt.

Evidence used: OpenAI: OpenAI Data Processing Addendum; reviewed May 21, 2026; Fresh.

Source evidence recorded: Official source linked: OpenAI Data Processing Addendum is linked as the official source for this sample finding.

Why this may matter: OpenAI provider-side retention answers are only part of the commitment. Product logs, traces, support tickets, and databases can retain the same prompts or outputs longer.

Unknowns: Whether debugging or observability tools capture customer prompts.

Owner decision: Security / GRC: Inventory where prompts, outputs, files, and embeddings are stored after the API call.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Anthropic: Tie Claude model-training answers to the product path actually used

Commitment: No specific customer commitment is mapped yet; treat this as a general review prompt.

Evidence used: Anthropic: Is my data used for model training?; reviewed May 21, 2026;

Fresh.

Source evidence recorded: Official source linked: Is my data used for model training? is linked as the official source for this sample finding.

Why this may matter: Claude API, Claude Enterprise, Claude Code, and unmanaged employee

use can require different evidence. A customer answer should name the Anthropic product path.

Unknowns: Whether employees use Claude outside the managed organization.

Owner decision: Security / GRC: Identify every Claude product used in product and employee workflows.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Anthropic: Verify whether zero data retention is actually covered by your Anthropic agreement

Commitment: Customer data is retained only as long as necessary.

Evidence used: Anthropic: Zero data retention agreement applicability; reviewed May 21, 2026; Fresh.

Source evidence recorded: Source note: Anthropic retention evidence needs agreement and product-scope validation before a zero-retention claim is reused.

Why this may matter: Zero data retention language should be tied to an approved agreement and the products it covers. It should not be used as a generic Claude statement.

Unknowns: Whether your team has a negotiated Anthropic retention agreement.

Owner decision: Legal / Privacy: Record agreement scope before using zero-retention language externally. Current handoff status: Needs agreement confirmation.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

- Azure OpenAI / Microsoft AI: Keep Azure OpenAI evidence separate from direct OpenAI evidence

Commitment: No specific customer commitment is mapped yet; treat this as a general review prompt.

Evidence used: Azure OpenAI / Microsoft AI: Data, privacy, and security for Models sold by Azure in Microsoft Foundry; reviewed May 21, 2026; Fresh.

Source evidence recorded: Official source linked: Data, privacy, and security for Models sold by Azure in Microsoft Foundry is linked as the official source for this sample finding.

Why this may matter: Azure-hosted model calls should cite Microsoft's Foundry data privacy source. Direct OpenAI API evidence should not be used for Azure deployments unless both paths are present.

Unknowns: Whether the selected model is covered by the Microsoft Foundry source.

Owner decision: Security / GRC: Confirm the Azure service, model, deployment region, and subscription.

Evidence limit: Review note only: the report does not provide legal advice, vendor certification, or final approval for customer answers.

Top findings

- OpenAI: Confirm OpenAI API data-use scope before reusing customer AI training language (P1, sample)

OpenAI API commitments should cite the platform data controls source and should not be copied to ChatGPT workspace or unmanaged account use without a separate review.

Source: <https://platform.openai.com/docs/guides/your-data>

- OpenAI: Separate provider retention from your own application retention (P1, sample)

OpenAI provider-side retention answers are only part of the commitment. Product logs, traces, support tickets, and databases can retain the same prompts or outputs longer.

Source: <https://openai.com/policies/data-processing-addendum/>

- Anthropic: Tie Claude model-training answers to the product path actually used (P1, sample)

Claude API, Claude Enterprise, Claude Code, and unmanaged employee use can require different evidence. A customer answer should name the Anthropic product path.

Source:

<https://privacy.anthropic.com/en/articles/7996868-i-want-to-opt-out-of-my-prompts-and-results-being-used-for-training-models>

Recommended actions

- Record the exact vendor products, plans, and data categories before reusing these answers with customers.
- Attach official source links and a last reviewed date to each Trust Center or questionnaire answer.
- Review DPA and subprocessor evidence before answering personal data questions.
- Add transfer-mechanism review to the follow-up list for EU personal data.
- Create a monitoring evidence packet for selected critical vendors.

Evidence gap handling

- Unavailable source checks stay visible: If a vendor source cannot be checked, the report names the affected document as a source coverage gap instead of treating it as clean evidence.
- Unknown applicability is preserved: When plan, region, agreement, model, or product-path scope is unclear, the report asks a review question rather than making a legal conclusion.
- Human approval remains explicit: AI Vendor Packet organizes official-source review evidence and suggested review actions. It does not provide legal advice or final approval for customer answers.

Limitations

- Sample findings are review prompts, not live detected changes, unless marked live.
- AI Vendor Packet does not provide legal advice or decide whether a vendor is compliant.
- Customer-facing commitments should be reviewed against the current vendor source and your actual implementation.