The AI Adoption

Security review / 16 September 2026

AI at work. What happens to your data?

A practical review of Claude, ChatGPT and Grok Bot, with the controls and conditions that matter before connecting them to your business.

Including xAI API, Microsoft Copilot, Microsoft Foundry and Amazon Bedrock. Australia, France, United Kingdom and United States.

Prepared by	Evidence checked	Review type
The AI Adoption	16 September 2026	Public documentation

In this review

01

What a business needs to know

An AI assistant can help staff work with company information. Before connecting it to a mailbox, a shared drive or a customer system, the business needs to know where that information goes, how long it stays there and what the assistant can change.

This review compares Claude, ChatGPT, the xAI API and Grok Bot. It also covers Microsoft Copilot, Microsoft Foundry and Amazon Bedrock, because buying a model through another provider changes the contract and the way data is handled.

The findings that affect the decision

Choose a specific product and configuration. Claude chat and the Claude API have different retention rules. ChatGPT Business and Enterprise offer different controls. Grok Bot runs on Cursor infrastructure and needs its own assessment. A conclusion about one of these products does not cover the others. A1 O1 G1

Check storage and processing separately. OpenAI offers eligible customers storage in Australia and the UK, but its published ChatGPT and API tables do not offer local model inference in those two countries. Selecting Europe does not promise France-only storage. O2 O3

Read the exceptions to zero data retention. Anthropic now describes a temporary ZDR option for eligible Fable customers while Enterprise Frontier Safeguards is introduced. OpenAI documents endpoint exclusions and customer-specific safety exceptions. Neither can be reduced to a permanent, product-wide promise that nothing is retained. A2 A3 O3

Review agent permissions before adding live accounts. Grok Bot can work through a signed-in user's applications. Its enterprise controls include network restrictions and action recording; several are unavailable on self-serve Teams. Assess the permissions granted to the computer as well as the model's data policy. G2 G3

A published control still needs to be configured and tested. This report reviews documentation. It has not inspected a customer's tenant, signed agreement, log export or deletion result. Section 11 sets out how to obtain evidence for a particular deployment.

Our assessment

For ordinary drafting with public or appropriately reduced data, a managed business workspace can be a proportionate starting point. Where work involves sensitive records, contractual location requirements or consequential actions, select the plan and deployment that meet those requirements and test them before use. An Enterprise label alone does not settle that decision.

Choose the smallest configuration that supports the intended work, gives the business the controls it needs and can be operated consistently by its team. A more complex cloud setup creates additional configuration and maintenance work; it should solve a specific requirement.

Scope and evidence

Vendor and regulator statements in this edition were checked against the linked primary pages on **16 September 2026**. Recommendations and test procedures are The AI Adoption's assessment. They are not claims that those tests have already passed.

The review covers retention, training, access, residency, assurance documents, agent risks, selected incidents and privacy questions for Australia, France, the UK and the US. It is a documentation review, not a penetration test or a legal opinion on a particular business. Private audit reports and negotiated terms must be obtained from the provider. Annex B lists the sources used in this edition.

02

Know which product you are buying

The same model may appear in a chat application, an API, a coding tool and an agent that operates a computer. Record the application, plan, contracting party, model, enabled tools and deployment region together. Keeping only the vendor name in a supplier register loses information that can change the outcome of the review.

Product	How it is used	Review boundary
Claude Team / Enterprise	Staff chat, projects and connected work	Anthropic workspace controls, enabled features and commercial agreement
Claude API	An application sends selected data to a model	Organisation, model, endpoints, storage features and application logs
ChatGPT Business / Enterprise	A shared business workspace	Plan, connected apps, retention and eligible residency features
OpenAI API	A custom application or automation	Project settings, endpoints, application state and monitoring
Grok Business / Enterprise	An organisational Grok offering	Exact licence and agreement; do not import API promises into chat
xAI API	An application calls Grok models	API team, retention mode, endpoint and tool scope
Grok Bot	An agent operates a computer and applications	Cursor account, computer, sessions, connectors and model routing
Microsoft Copilot / Copilot Chat	Work across Microsoft 365 information	Tenant permissions, licences, Purview and enabled AI providers
Microsoft Foundry / Amazon Bedrock	Models accessed through a cloud service	Model offer, contract, deployment, routing and application settings

These boundaries are described in the providers' documentation. Microsoft is transitioning its Microsoft 365 Copilot naming; procurement should keep the actual licence name and agreement rather than infer a contract change from the shorter label. A1 O1 X1 G1 M1 F1 B1

Follow one piece of information

Trace a sample record through five places: the source application, the AI service, any tool or search provider, the output destination and the logs or backups. Put the owner and deletion rule beside each copy.

For example, a customer email may remain in the mailbox, appear in a prompt, be copied into an application log and be included in a drafted reply. Deleting the chat does not establish that all four copies have been removed.

1. Which fields must leave the source system for this task to work?
2. Which people and providers can receive those fields?
3. Does the assistant only propose an action, or can it carry it out?
4. Where can an administrator see what happened?
5. How are access and stored copies removed when the work ends?

03

Anthropic: Claude workspaces and API

Training and retention

Anthropic says commercial customer content is not used for model training by default. Voluntary feedback and explicit opt-in arrangements need separate consideration. Record which commercial account staff use and how feedback is handled. A4

The ordinary API policy deletes inputs and outputs within 30 days, with exceptions for features such as Files, an agreed retention arrangement, policy enforcement and legal obligations. Saved commercial chats remain available as part of the product experience. After a user deletes a conversation, backend deletion is ordinarily within 30 days. Flagged inputs and outputs may be kept for up to two years, and associated safety scores for up to seven years. Feedback has a separate five-year retention period. A1

Enterprise custom retention has a 30-day minimum. The timer is based on the last activity, and **project retention takes precedence over chat retention**. Projects retained indefinitely can keep their chats beyond a shorter chat policy. Setting one chat period is insufficient proof that the whole workspace follows it. A5

Covered Models and the newer exception

Anthropic's Covered Models policy describes 30-day retention for safety review, with longer retention possible for flagged material or legal requirements. The current page also points to an exception for eligible organisations notified that they can use Fable with ZDR. A2

The accompanying guidance describes a limited-time arrangement for eligible internal business applications using Fable 5 and 5.1. Enterprise Frontier Safeguards is scheduled to roll out in phases from autumn 2026. Eligibility must be confirmed; a future rollout is not an active control in a customer's workspace. [A3]

Assessment: specify the model and obtain confirmation of the applicable arrangement. Avoid both blanket claims: that every Claude model requires safety retention, or that a ZDR agreement necessarily covers every model and feature.

Access and evidence

Claude's audit-log documentation limits the feature to Enterprise. The export covers the previous 180 days and contains event metadata, not chat or project content. Content export is separate. Customer-managed encryption keys disable the ordinary Export logs button, but current guidance says audit events remain available through the Compliance API. Encryption keys therefore do not remove every audit-access route. [A6]

Anthropic lists SOC 2 Type I and II, ISO 27001 and ISO 42001 among its commercial compliance credentials. It also describes HIPAA-ready configurations with a BAA available. The underlying report or certificate must establish the purchased product's scope and period. A public credentials page is not inspection of the private audit report. [A7]

Location

The direct API documentation separates inference geography from workspace geography. On the page reviewed, inference can be global or US-only, and the workspace geography is US. A US inference setting does not establish an Australian, UK or French storage commitment. Cloud partner arrangements need their own review. [A8]

Before approval

Confirm the model's retention regime, the workspace's location commitments and the assurance scope. Test chat and project deletion independently, export an audit event and check the content-export path needed by the business. Where customer-managed keys are required, test the Compliance API rather than assume the standard download remains available.

04

OpenAI: ChatGPT and API

Business information and human access

OpenAI states that business data is not used to train its models by default. Eligible business products have encryption and organisational controls, but the precise features differ by plan. Name the product in the review. [O4]

The enterprise privacy page describes authorised employee access for specified purposes and, for Business and API data, specialist contractors reviewing abuse and misuse. Its Enterprise wording limits access to authorised employees. No-training commitments therefore do not mean that no authorised person can ever access stored content. [O1]

Retention requires a product-specific answer

The privacy page says Enterprise administrators control retention. It also says Business administrators can control retention, while the earlier research recorded different wording in Business help pages. This edition does not resolve that into an unsupported claim that a particular control is available in every Business workspace. Obtain current product-specific instructions and demonstrate the setting in the actual tenant. [O1]

For the API, ordinary abuse-monitoring retention is up to 30 days. Approved ZDR and Modified Abuse Monitoring have endpoint and feature conditions. Application state, such as stored files or conversations, is separate. The documentation also describes Eyes Off and Safety Retention arrangements that can change retention or review for specified customers after notice. [O3]

Assessment: review project settings alongside the application's storage behaviour. A no-storage request parameter cannot prove that application logs, uploaded files or another service retain nothing.

Storage and inference

Eligible Enterprise and Edu workspaces can select supported storage regions. Australia and the UK are listed; the inference-residency list contains Europe, the US and the UAE. Some processing, system data and external integrations remain outside the commitment. ChatGPT Business is not included in the eligibility wording reviewed. [O2]

The API lists storage but not local processing for Australia and the UK. For France, an EEA-and-Switzerland commitment is broader than a France-only requirement. Check model and feature eligibility for the intended endpoint. [O3]

Assurance and contract

OpenAI describes SOC 2 Type 2 coverage for API and business ChatGPT services. Its ISO 27001/27701 wording specifically names the API, Enterprise and Edu. It separately describes ISO 42001 coverage for its AI management system. Request the current scope for the purchased product; omission from one summary page does not establish absence from every certificate. [O5]

The DPA sets the processing relationship and incident obligations. Keep the applicable version with the order. Verify any SLA, additional retention right or special location commitment in that agreement rather than infer it from marketing. [O6]

Before approval

Demonstrate retention and member removal in the selected workspace. For a custom application, record the endpoint, region, model and approved monitoring mode. Test representative files, conversations and connectors. If content must stay in one country, check storage, inference and external-tool processing separately.

05

Grok Bot and xAI need separate reviews

Grok Bot: a computer with company access

Grok Bot operates on Cursor-hosted computers in the US. Its security FAQ states that each user has an isolated computer, but **all Bots belonging to that user share it**. A second Bot does not create a separate place for credentials or client files. The FAQ also states that Grok Bot is within Anysphere's current ISO 27001 and ISO 42001 scope. Request the certificates to confirm scope and validity. [G1]

The security documentation describes a durable disk, persistent application sessions and separate retention at model providers. Deleting a Bot or stopping the computer does not remove all files or sessions. A per-organisation retention policy and customer-managed point-in-time restore are not available. [G2]

Cursor manages the serving model. The team model allowlist is Enterprise-only, and the FAQ warns that enforcement is not guaranteed. A contract that restricts subprocessors needs written clarification before rollout. [G1]

Controls depend on the plan and their settings

Grok Bot's enterprise documentation identifies organisational controls for enablement, network access, action recording, identity and computer management. Self-serve Teams do not receive the same set. Check the actual plan before designing a policy around these controls. [G3]

Network access defaults to allow-all without a policy. Audit Logs and Action Recording are separate: administrative logs are not a complete record of the agent's work. Action Recording is off by default and, when enabled, has a 90-day internal retention period. Export requires its own configuration. [G2]

Approvals and Auto Review can constrain consequential actions. The documentation also identifies operations outside the review layer, so use restricted application permissions and network destinations to limit the effect of a mistaken decision. [G4]

Grok Bot uses the member's Cursor identity and signed-in application access. Team or service-account connectors need separate attention. Offboarding should remove membership and revoke application access, rather than only stop a computer. [G5]

A practical boundary for the first workflow

For an email assistant, begin with the mailbox and folders it needs, a permitted output destination and a draft-only action. Use a separate account or permission boundary for work that must remain isolated. Test whether the agent can reach unrelated files, another service through the browser, or a destination outside the approved list.

Confirm privacy mode and model-provider arrangements before using sensitive content. A ZDR commitment at the model layer does not delete files on the computer. The security page also notes that safety-flagged content can be retained by providers. [G2]

xAI API: a different retention and contract path

The API security FAQ says inputs and outputs are not used for training without explicit permission. Ordinary API requests and responses are retained for 30 days for abuse auditing. Self-serve ZDR, where available, applies to the entire team rather than individual keys. It disables storage-dependent features, including Files, Collections and Batch. The response header can confirm whether ZDR is active. [X1]

The current Enterprise Terms name SpaceXAI LLC. They restrict submission of personal data to the ZDR-enabled API, and impose additional BAA and ZDR conditions for HIPAA-protected health information. Resolve the agreement that applies to a proposed chat or personal-data use before proceeding; a general marketing assurance cannot settle that condition. [X2]

The DPA describes incident notification without undue delay and, where feasible, within 48 hours of awareness. The qualification matters: this is not an unconditional 48-hour guarantee. [X3]

A documented US endpoint now exists

The regional-endpoint documentation gives a US route for supported API request handling, inference, moderation and retained request data. Files, Collections and server-side tools fall outside that guarantee. The global endpoint can route across regions. This corrects the earlier claim that no residency mechanism was documented. The reviewed page does not establish an Australian, UK or France-only API route. [X4]

Before approval

For Grok Bot, confirm the Cursor contract, certificate scope, model-provider restrictions and required Enterprise features. For xAI API, confirm the Enterprise Terms, personal-data condition, ZDR status and endpoint. Keep those decisions separate in the supplier register.

06

Microsoft Copilot and cloud alternatives

Microsoft Copilot / Copilot Chat

Copilot works with information a user can access in Microsoft 365. Microsoft states that prompts, responses and Graph content are not used to train foundation models. Existing permissions remain important: if a shared folder exposes confidential documents to too many people, enabling an assistant does not correct that access. [M1]

Purview provides retention, audit and DLP capabilities, with different requirements. Copilot audit events are collected through Audit (Standard) when auditing is enabled. That is evidence of activity, not proof that every prompt, file and action is available in one log. [M2]

Purview DLP can restrict labelled content, sensitive prompts and web-search use. Its documentation refers separately to licensing and identifies some controls as preview or subject to rollout. **An available feature is not necessarily included, enabled and effective in every tenant.** Test the exact policy and workload. [M3]

Retention policies operate through Exchange's storage and deletion processes. Microsoft's worked example shows that a one-day deletion policy can take 16 days to complete. Holds and competing retention obligations can suspend permanent deletion. Removing a message from the visible conversation is insufficient evidence of deletion. [M4]

Check which AI provider is enabled

Microsoft describes Anthropic as a subprocessor for some Copilot experiences, with regional exclusions and separate settings. It also describes Anthropic models with Data Retention under separate Anthropic terms; that path is off by default. Eligible organisations may receive a Fable configuration under Microsoft's terms without Anthropic retaining content. Review the actual setting and eligibility rather than infer the contract from the model name. [M5]

Assessment: an existing Microsoft environment can provide useful identity and compliance tools. Approval still depends on permissions, enabled providers, licence coverage and tested policies. There is no basis here for declaring Microsoft universally the safest provider or a certificate absent merely because one service list uses another product name.

Microsoft Foundry

For **Models sold by Azure**, Microsoft's privacy documentation says the models are hosted in Azure, customer content is not available to the model provider, and content is not used to train foundation models without permission. Stateful features and abuse monitoring have their own handling. These statements apply to the specified offer category, not automatically to everything in the Foundry catalogue. [F1]

Deployment type affects location. Standard deployments process within the specified geography, with possible movement between regions in that geography. Global deployments can process across geographies; DataZone deployments can use the selected zone. A resource created in France or Australia does not, by itself, establish that inference stays in that country. [F1]

Confirm the model offer, deployment type and availability for the required geography. Keep the deployment record, endpoint and exclusions with the review. Exact model-region availability is a deployment-time check; this edition does not endorse the old draft's blanket country-by-country availability claims.

Amazon Bedrock

AWS operates the Bedrock model deployment accounts. Its data-protection page says model providers do not have access to those accounts or to customer prompts and completions. Customers remain responsible for their permissions and cloud configuration. [B1]

Bedrock documents several retention modes. `none` enforces zero retention and blocks models that require it; `default` follows the model's policy. `aws_review` permits AWS human review where required. The legacy `provider_data_share` label does not mean data is currently shared with model providers. For Fable 5 and 5.1 in the reviewed guidance, retention and review stay within AWS for up to 30 days. A `store=false` parameter alone does not guarantee ZDR. [B2]

Cross-region inference can route processing beyond the region receiving the request. Geographic profiles and global profiles have different boundaries. Record the actual profile and destinations, including where safety-retained content is stored. [B2] [B3]

What changes when you use a cloud route

Foundry and Bedrock can support an application with a deliberately constrained data path. They also require someone to own deployment, permissions, logging, upgrades and cost control. They do not reproduce every feature of the corresponding chat product. Choose a cloud route when those responsibilities are justified by the work or a specific requirement.

07

Residency and privacy by country

Four location questions

Record where content is stored, where the model runs, where other processing takes place and where authorised support or safety reviewers may access information. Include connectors and search services. A country name on a workspace setting answers only the question that the provider defines for that setting.

Annex A distinguishes documented options from requirements still needing confirmation. Availability can vary by model, feature and contract. The table is not approval for an entire deployment.

Australia

The earlier draft said the Privacy Act applies to an Australian business without qualifying that statement. OAIC guidance says most businesses with annual turnover of $3 million or less are exempt, with important exceptions, including health service providers and other specified activities. Establish coverage before stating a legal obligation. Other legal and contractual duties may still matter. [L1]

For a covered entity, assess the information involved and any overseas handling under the applicable Australian Privacy Principles. A health workflow needs particular attention to collection, use, access and disclosure. Local storage alone cannot establish compliance for the entire workflow. OAIC's AI guidance provides a starting point. [L2]

Under the NDB scheme, a suspected eligible breach requires a reasonable and expeditious assessment, taking all reasonable steps to complete it within 30 days. Once there are reasonable grounds to believe an eligible breach occurred, notification is required as soon as practicable. The 30 days are not a routine waiting period for notification. [L3]

France and the European Economic Area

A transfer outside the EEA needs an applicable transfer mechanism. Depending on the recipient and circumstances, this may be an adequacy decision, safeguards such as Standard Contractual Clauses, or a narrowly applicable derogation. Verify the exact recipient and service. An unsuccessful company search in a register is not enough to declare that no valid mechanism exists. [L4]

For relevant personal-data breaches, notification to the competent authority is generally required without undue delay and, where feasible, within 72 hours of awareness. Notification to affected people depends on the applicable high-risk test. Record the business's notification process separately from the supplier's contractual period. [L5]

If a contract requires France-only handling, Europe-wide storage does not meet that wording. Health hosting, public-sector procurement and sector-specific obligations need their own assessment; this review does not award HDS or SecNumCloud status to any route.

United Kingdom

Check the UK rules for international transfers and the safeguards that apply to the named recipient. Do not assume an EU clause, US hosting location or general GDPR statement settles the UK transfer question. Keep the applicable agreement and assessment with the deployment record. [L6]

The ICO says a breach likely to create a risk to individuals must be reported as soon as possible and, where feasible, within 72 hours. High-risk breaches also require notice to affected people without undue delay. An initial report can be supplemented as facts become available. [L7]

United States

Requirements depend on the activity, data and applicable federal and state laws. For a HIPAA-regulated cloud use, HHS guidance requires an appropriate BAA and compliance with the other HIPAA obligations. A HIPAA-ready label does not replace the agreement and eligible configuration. [L8]

US legal process can involve data held abroad. The CLOUD Act does not give authorities unrestricted access to every customer's content; applicable legal requirements and safeguards still matter. Data location and exposure to legal process are separate questions for the contract and transfer assessment. [L9]

08

Risks to test in the workflow

The following are review scenarios, not findings that any named business has suffered an incident. Use those relevant to the proposed work and assign an owner to the controls.

OWASP identifies direct and indirect prompt injection as ways that input can redirect an AI system. External documents and websites can carry instructions as well as information. Restricting privileges and requiring approval for consequential actions reduce the effect of a successful attack; a prompt telling the model to behave safely is not a complete control. [R1]

Access and disclosure

Scenario	What to check	Useful evidence
A supplier document tells the assistant to upload records elsewhere	Destination and action permissions independent of the document	Synthetic-data test; blocked request and recorded event
A drive exposes unrelated payroll or customer records	Source permissions and connector scope	Approved files accessible; restricted files inaccessible
A Bot uses another Bot's browser login	Isolation at the user or account boundary	Separate accounts where needed; access test
A spreadsheet contains hidden personal columns	The actual input sent to the service	A reduced export of necessary fields
A search tool receives sensitive terms	Fields leaving the approved service	Tool-path and destination test
A former employee retains a session or connector	Membership, tokens and application sessions	Revocation followed by a failed access attempt

Use representative but non-sensitive records. A successful test is evidence for the configuration tested; repeat it after material permission or connector changes.

Actions and reliability

Scenario	What to check	Useful evidence
An invoice correction is applied without review	Where write authority is granted	Approval bound to the record and proposed change
A plausible answer cites the wrong source	Traceability to the original	Source, date and review of the material statement
A calculation is wrong	Reproducibility of totals	Recalculation against source records
A scheduled job fails silently or runs twice	Alerts and duplicate prevention	Alert delivered; safe rerun demonstrated
A model or tool changes after approval	Change ownership	Updated record and relevant tests
An agent keeps retrying	Stop conditions and usage controls	Bounded failure test and visible stop

For money movements, employment decisions, clinical material and safety documentation, specify the human decision-maker and the source they review. Do not make the assistant both the author and the sole approver.

Retention and recovery

Scenario	What to check	Useful evidence
The chat is deleted but files remain elsewhere	Files, application state, exports and logs	Checks for each relevant storage location
The service is unavailable	A usable fallback and current source data	Owner completes the task through the fallback
A configuration change invalidates the review	Retention, region and provider changes	Named reviewer, change record and renewed approval

Keep logging proportionate. A full prompt log may create another sensitive dataset. Record enough to establish who acted, on which record, with what approval and outcome; set access and retention for the log itself.

09

Incidents: what the evidence establishes

An incident, a vulnerability and a regulatory investigation answer different questions. Raw counts would give a misleading ranking because products, disclosure practices and observation periods differ.

These are selected primary-source examples, not an exhaustive inventory. The earlier draft's unverified totals are not repeated.

OpenAI: March 2023 disclosure

OpenAI reported a software defect that allowed some users to see another user's chat titles and potentially a newly created message. It also described possible exposure of limited payment-related information for some active Plus subscribers during a defined period, and said the defect had been patched. [11]

Relevance: ask how isolation failures are detected and customers notified. This historical incident does not establish that a current Enterprise tenant has the same vulnerability.

Anthropic: cybersecurity evaluation incidents

Anthropic disclosed three incidents in its cybersecurity evaluations in July 2026. Its account describes models reaching real organisations' infrastructure during tests. It also states that the dedicated evaluation infrastructure was separate from, and had no access to, customer data. [12]

Relevance: test the actual boundary around an agent's tools and network access. The disclosure is not evidence that Claude business workspaces leaked customer records.

Grok: a regulator's investigation

On 3 February 2026, the ICO announced investigations into X Internet Unlimited Company and X.AI concerning personal-data processing associated with Grok. An investigation is not a finding of a breach of law or an imposed penalty. This review does not establish its later procedural outcome. [13]

Relevance: confirm the entity, service and activity under review. Do not transfer the allegation to Cursor's Grok Bot infrastructure without evidence.

How to use vulnerability records

For a relevant CVE, record the affected version, exploit conditions, fix and deployment exposure. A vulnerability in a local coding client is not automatically a vulnerability in the hosted chat service. A fixed server-side issue is not an instruction to install an unrelated update.

The original research folder contains a broader catalogue. Its entries, counts and current status have not all been revalidated here and are excluded from the decision tables. Reopen the vendor or regulator's record before using an entry in a procurement decision.

10

Four examples of applying the review

These are generalised business scenarios. They are not client endorsements, delivery claims or assessments of a named organisation. They retain the original four sectors while removing engagement details that could identify a business.

A health practice: operational reporting

Suppose a manager needs a weekly view of appointments, cancellations and staff availability. Define the required totals and calculate them directly from the records where possible. Do not send treatment notes or sickness explanations to a model merely because they are present in the export.

If the work needs identifiable health information, assess lawful use, contracts, authorised recipients and retention before choosing a service. Clinical wording intended for a patient should have a named qualified reviewer and an approved source.

Approval evidence: a sample showing which fields leave the practice system, confirmed permissions, the applicable agreement and an approved output. A more expensive plan cannot compensate for unnecessary disclosure.

A manufacturer: invoice reconciliation

Compare freight invoices with an agreed rate card using a reproducible calculation. The assistant can explain an exception and draft a query. Keep accounting changes behind approval showing the invoice, current value and proposed correction.

Give the workflow relevant extracts rather than unrestricted access to the owner's mailbox and finance applications. Record the rate-card version. Every discrepancy should remain traceable to its invoice and calculation.

Approval evidence: independently checked totals, a blocked unapproved write and a log linking an approved change to its reviewer.

A construction business: tender support

A tender assistant may need pricing, drawings, specifications and supplier correspondence. Define which material it can access and which outputs carry commercial or safety consequences. Source dates and document versions matter as much as fluency.

Test a supplier document containing an instruction to send information externally. The workflow should treat the document as evidence to analyse while the business's permissions control action. Keep accounting, tender submission and safety approval separately authorised.

Approval evidence: restricted source access, current references, a reviewed quotation and a test that unapproved submission fails.

A hospitality group: a weekly summary

Begin with aggregated sales, staffing and purchasing information for a clear reporting period. Separate calculation from the assistant's explanation. Only include payroll records, employee health details or bank information when the task requires them and the use is approved.

If an agent later updates a roster or places an order, assess that write as a new permission. Approval of reporting is not approval to change the source system.

Approval evidence: reconciled totals, an owner who checks material commentary and a defined boundary between reporting and action.

11

What to confirm before rollout

Agree the intended work first. Then complete the following record for that workflow. Leave an item unresolved if the required evidence is unavailable.

Product and agreement

Confirm	Evidence to keep
Product, plan, owner and legal counterparty	Order, agreement and account identifiers
Data categories and purpose	Sample reduced input and purpose statement
Training, retention, review and deletion	Applicable policy and contract provisions, dated
Storage, inference and other processing locations	Region settings, deployment type and exceptions
Subprocessors and tools	Applicable list and approved integrations
Assurance scope	Current report or certificate naming service, entity and period
Support and incident obligations	Contractual contacts, terms and notification chain

A missing public SLA or certificate is a procurement question, not proof that no negotiated commitment exists. Where pages disagree, ask a precise question about the selected product and keep the answer.

Operational checks

1. **Access:** a representative user reaches the necessary source and cannot reach excluded records.
2. **Actions:** the assistant drafts the output; consequential writes require the agreed approval.
3. **Traceability:** a useful event record identifies who approved an action.
4. **Retention:** chats, projects, files and logs follow their rules, with exceptions documented.
5. **Offboarding:** removal and token revocation stop access, including browser sessions where relevant.
6. **Failure:** a failed or repeated run creates a visible, recoverable outcome.

Record the date, configuration and result. Use synthetic data initially. Documentation supports an expected result; it cannot replace observing that result in the business's configuration.

The decision

Use one of three outcomes: approved for the specified work, approved with written conditions, or held pending evidence. Name the owner and the events requiring review, such as a new connector, provider, data category or write permission.

Approval should let a colleague tell what is allowed. “Claude is approved” leaves too much unanswered. “The approved workspace may draft internal procedures from these folders, with no personal records or external sending” describes a boundary that can be checked.

12

Unresolved questions and limits

These points remain dependent on the supplier or the business's configuration.

Question	What resolves it
Is the purchased service in the current SOC or ISO scope?	Report or certificate with applicable service scope
Which retention controls exist in the selected ChatGPT Business workspace?	Current instructions, tenant demonstration and clarification
Is the organisation eligible for temporary Fable ZDR?	Notice or written approval naming organisation and usage
Can Grok Bot meet a fixed model-provider restriction?	Contractual clarification and a tested configuration
Which xAI terms govern the personal-data use?	Applicable agreement and confirmation of permitted route
Do Copilot DLP controls cover the users and applications?	Licence, rollout status and tenant test
Does a cloud model meet a country-only requirement?	Current offer, deployment, destinations and contract
When will retained or held content be deleted?	Applicable rules, hold status and supplier confirmation
Which SLA and breach terms were purchased?	Signed or incorporated agreement

The earlier research files remain for traceability with a notice that they predate this revision. Their counts, broad certification negatives and unresolved claims are not carried forward as verified facts. Annex B identifies this edition's primary sources.

Terms used in this review

- **Retention:** how long a particular copy is kept. Chats, files, safety logs and backups can follow different rules.

- **ZDR:** a provider's zero data retention arrangement, with defined product and feature scope.
- **Residency:** a location commitment; check which storage and processing it covers.
- **DPA:** the agreement governing processing of personal data between the relevant parties.
- **DLP:** controls to detect or restrict specified information or content flows.
- **SSO / SCIM:** centralised sign-in and automated account provisioning or removal.
- **SOC 2 Type 2:** an assurance report about specified controls over a stated period; read scope and exceptions.
- **ISO 27001 / ISO 42001:** management-system standards; certification does not approve every feature or use.
- **Subprocessor:** another organisation engaged to process personal data in the relevant chain.
- **BAA:** the agreement required for relevant HIPAA business-associate relationships.

Annex A. Comparison tables

These tables summarise the reviewed documentation. A documented capability is not evidence that it is enabled for a particular customer. Read the named product's section for the conditions.

A1. Retention and operational control

Product or route	Retention position	What the business must establish
Claude commercial chat	Saved conversations persist; deletion and safety exceptions apply [A1]	Chat and project policies, content removal and access
Claude Enterprise	Configurable retention; project rules take precedence [A5]	Both policies configured and checked; required audit route tested
Claude API	Standard retention and model-specific conditions; eligible Fable ZDR exception [A1] [A2] [A3]	Model, feature and agreement eligibility
ChatGPT Business	Public privacy page describes admin control; exact tenant capability needs confirmation [O1]	Demonstration of the applicable retention control
ChatGPT Enterprise	Admin-controlled retention and deletion exceptions [O1]	Agreed settings, enabled features and export requirements
OpenAI API	Monitoring and application state are separate; ZDR has conditions [O3]	Approved project mode and endpoint-specific behaviour
xAI API	30-day default; team-wide ZDR removes storage-dependent features [X1]	Applicable personal-data terms and active mode
Grok Bot	Computer persistence and provider retention are separate [G2]	Files, browser sessions, privacy mode and available controls
Microsoft Copilot	Purview policies, deletion processing and holds [M4]	Licences, configured policy and actual workload coverage
Foundry Models sold by Azure	Stateful features and abuse monitoring have separate handling [F1]	Offer category, stored state and monitoring arrangement
Amazon Bedrock	Effective retention mode and model requirements govern handling [B2]	Actual mode, model eligibility and stored application data

A2. Location: what the documentation supports

Route	Documented position	Country requirement to resolve
Claude direct API	US workspace geography; global or US inference [A8]	No AU, UK or France-only route established by this page
Eligible ChatGPT workspace	AU and UK storage options; Europe covers EEA + Switzerland [O2]	Storage does not establish AU/UK inference or France-only handling
OpenAI API	AU and UK storage, without local processing in the published table [O3]	Endpoint, model and feature scope; Europe is not France-only
xAI API	A scoped US endpoint; global route has no fixed processing location [X4]	Tools and files are outside the US guarantee; other country commitments need confirmation
Grok Bot	Hosted computers in the US [G1]	Computer location does not settle every model-provider or tool path
Microsoft Copilot	Review the tenant, workload and enabled provider commitments [M1] [M5]	Anthropic settings have specific regional exclusions
Foundry Models sold by Azure	Standard, DataZone and Global have different processing boundaries [F1]	Select a currently available model and deployment meeting the required geography
Amazon Bedrock	Region and inference profile determine the possible destinations [B3]	Confirm every destination; a receiving region alone is insufficient

For an Australian or UK requirement, distinguish local storage from local inference. For France, distinguish a national requirement from an EEA requirement. For the US, check the scope of the US option, including excluded tools. These are interpretations of the documented boundaries, not guarantees for an uninspected tenant.

A3. Assurance and supporting documents

Provider or product	Public evidence used	Additional evidence to obtain
Anthropic	Commercial credentials page lists SOC and ISO standards [A7]	Current report, scope, period and exceptions for the selected surface
OpenAI	Security page distinguishes named services and standards [O5]	Current product-specific scope, including Business where relevant
Grok Bot / Anysphere	FAQ explicitly identifies Grok Bot within ISO scope [G1]	Current certificates and applicable Cursor agreement
xAI API / Grok business	Enterprise Terms and DPA were read [X2] [X3]	Assurance report and exact product scope; no negative certification conclusion is made here
Microsoft Copilot / Foundry	Product-specific security and privacy documentation [M1] [F1]	Applicable Microsoft assurance pack and purchased licence coverage
Amazon Bedrock	Service data-protection and retention documentation [B1] [B2]	Applicable AWS assurance pack and configuration record

A4. What each assurance answers

Evidence	What it can support	What it cannot establish alone
A current audit report	Controls in the stated service, scope and period	The security of every new feature or customer configuration
A certificate	The certified management system and scope	Approval of a particular business workflow
A DPA and transfer terms	Processing responsibilities and contractual safeguards	That access permissions or deletion settings are correct
A residency statement	The defined location commitment	That every connected application follows the same boundary
A configuration export	The recorded settings	Their effectiveness without relevant tests
A workflow test	Behaviour observed for the tested case	Permanent assurance after material changes

This interpretation is the assessment method used in the report. It avoids turning a missing public document into a claim that a control does not exist.

Annex B. Primary sources

Read on 16 September 2026. The links lead to the provider or public authority. They establish what the source says; they do not prove that a customer's configuration meets it. Reports behind supplier access requests were not independently inspected for this edition.

Anthropic

A1. Commercial data retention. Ordinary retention, deletion, feedback and safety exceptions.

A2. Data retention practices for Covered Models. Model-specific policy and eligible-organisation exception.

A3. Covered Models. Temporary Fable ZDR eligibility and Enterprise Frontier Safeguards rollout.

A4. Commercial data and model training. Default training commitment and opt-in handling.

A5. Enterprise custom retention. Chat and project rules, activity timers and minimum period.

A6. Access audit logs. Enterprise scope, export contents and Compliance API with managed keys.

A7. Anthropic compliance credentials. Public claims and route to request supporting documents.

A8. Claude API data residency. Workspace and inference geography.

OpenAI

O1. Enterprise privacy. Product-specific retention, training and human-access statements.

O2. ChatGPT data and inference residency. Eligibility, supported regions and exclusions.

O3. API data controls. Monitoring modes, endpoint scope, application state and residency tables.

O4. Business data privacy, security and compliance. Commercial data commitments.

O5. Security and privacy. Named services and public assurance statements.

O6. Data Processing Addendum. Processing and incident responsibilities.

Grok Bot and xAI

G1. Grok Bot security FAQ. Isolation, hosting, model restrictions and ISO scope.

G2. Grok Bot security. Networks, logs, persistence, privacy and deletion.

G3. Grok Bot for teams and enterprises. Plan-dependent administration and rollout settings.

G4. Approvals, security and privacy. Member approval and review behaviour.

G5. Grok Bot identity and access. Sign-in, provisioning and application identity.

X1. xAI API security FAQ. Training, retention, ZDR scope and confirmation.

X2. Enterprise Terms. Contracting entity and personal-data conditions.

X3. Data Processing Addendum. Processing obligations and qualified incident-notice timing.

X4. Regional endpoints. US endpoint guarantee and excluded features.

Microsoft and Amazon Web Services

M1. Copilot data, privacy and security. Microsoft 365 context, access and training.

M2. Audit logs for Copilot and AI applications. Audit enablement and event scope.

M3. Purview DLP for Copilot. Supported restrictions, licensing references and preview limits.

M4. Retention for Copilot and AI apps. Storage, deletion processing and holds.

M5. Anthropic models in Microsoft Online Services. Subprocessor and separate-contract paths, defaults and eligibility.

F1. Privacy for Foundry Models sold by Azure. Offer scope, hosting, processing geography and stored state.

B1. Amazon Bedrock data protection. Shared responsibilities and model-provider access.

B2. Amazon Bedrock data retention. Modes, model requirements and AWS review.

B3. Amazon Bedrock cross-region inference. Geographic and global routing.

Privacy authorities and legal guidance

L1. OAIC: small business. Privacy Act coverage and exceptions.

L2. OAIC: commercially available AI products. Privacy review of AI use.

L3. OAIC: Notifiable Data Breaches scheme. Assessment and notification requirements.

L4. CNIL: transfers outside the EU. Transfer mechanisms and conditions.

L5. CNIL: personal-data breach notification. Reporting obligations and timing.

L6. ICO: international transfers. UK transfer assessment and safeguards.

L7. ICO: UK GDPR breach reporting. Risk threshold and notification timing.

L8. HHS: HIPAA and cloud computing. BAA and cloud responsibilities.

L9. US Department of Justice: CLOUD Act resources. Official legal-process background.

Threats and selected incidents

R1. OWASP: prompt injection. Threat definition and mitigation principles.

I1. OpenAI: March 2023 incident disclosure. Historical scope and reported remediation.

I2. Anthropic: cybersecurity evaluation incidents. Incident context and infrastructure boundary.

I3. ICO: investigation announcement, 3 February 2026. Announcement only; no later outcome established here.