Enterprise Data Security & Privacy

A guide for organisations evaluating Claude Code for workforce use, in two layers: first the whole answer for the leadership team, in a page; then the depth an IT, security and compliance review needs — what leaves the machine, where it goes, how long it exists, who is legally responsible for it, and what you can control centrally — with links to the primary sources throughout.

THE ANSWER FOR LEADERSHIP

This section is for the person deciding. Everything after it is for your security team — forward the whole document.

Where does your data go? Nowhere without your say-so — that is the design.

Your organisation's working files stay on your people's own computers, as plain files you can open and read. Nothing is uploaded in bulk, no index of your data is built anywhere, and nothing runs in the background. Data moves only when a person working chooses to send it: task by task, they direct which files the AI reads and supervise everything that comes back.

That discipline is not an add-on to the training — **it is the training**. Managing what the AI reads for each piece of work is the core skill we teach, so the same practice that makes your people effective with AI is what keeps your data governed.

Your organisation holds its own layer of control on top of that: policies for what may be shared, deny rules that put named files beyond the AI's reach, central settings no individual can override — and the commercial agreements governing the AI itself, under which the model provider is **contractually prohibited from training on your data**. That is the default, not something you opt into.

The rest of this document takes your security team through the full picture, with primary sources throughout. They will be in good company: the sectors with the strictest confidentiality duties — major banks, most of the UK's top law firms, healthcare, the US federal government — have completed this review and deployed.

One pattern worth carrying into the review: organisations whose biggest barrier to AI is data control tend to discover that this is the version of AI they can actually approve. **The design that makes it governable is the same design that makes it work.**

THE SECURITY REVIEW

How much of this document applies to you

For most organisations, adoption comes down to one decision: run Claude on a commercial plan (Team or above). Every protection that matters in this document — the contractual no-training prohibition, 30-day deletion, processor status under a Data Processing Addendum, encryption in transit and at rest — comes with that plan by default. Nothing has to be configured, negotiated or deployed to get it. A five-minute managed-settings baseline (deny rules on credential files, telemetry off) is worth adding and is described below, but it is a hardening step, not a precondition.

Everything beyond that — zero data retention, cloud-perimeter deployment, the Compliance API, US-only inference — exists for organisations whose data classification demands it. If yours does, your security team will already know; most are not in that position, and none of it needs deciding before people start.

Individuals exploring before their organisation commits can do so on a personal plan with the training setting turned off — the first module of the training walks through that setting.

© Multiply Academy 2026 · multiply.academy

The review in three sentences

> On commercial plans, Anthropic is **contractually prohibited from training on your data** and acts as your data processor. Claude Code sessions are encrypted in transit and at rest and deleted from Anthropic's servers within 30 days — with zero-retention arrangements available for qualified enterprise accounts. Every control can be enforced centrally through managed settings, including disabling all non-essential transmission.

A deliberately narrow scope

Multiply Academy's training deploys one AI tool: **Anthropic's Claude Code**, used interactively in a terminal on the employee's own computer. There is no web workspace to govern, no tenant integration, and no background service. The entire data story is a single, inspectable loop — which makes this a materially simpler review than most AI tools your team will have assessed.

If your organisation has already approved Microsoft 365 Copilot — the most common prior AI review — the questions here have the same class of answers (commercial no-training terms, encryption, central governance), applied to a much narrower surface: Copilot's review concerned tenant-wide access to mail, calendars and SharePoint by design; Claude Code starts with access to nothing beyond the files a user's task touches, and any reach into other systems exists only if deliberately connected — per user, under that user's own sign-in, governable centrally (see *Connecting other systems* below). The one genuinely different consideration is that processing happens on Anthropic's infrastructure rather than inside your Microsoft tenant — the retention and control sections below address it directly.

THE DATA FLOW IN ONE PARAGRAPH

An employee's files stay on their computer. When they set a task, Claude Code sends the instruction plus the specific files it reads for that task — encrypted — to Anthropic's servers; the model's response comes back; the loop repeats. The person working directs what is read and reviews what comes back — and because choosing what the AI reads per task is the core skill the training teaches, that human control point is a practised discipline, not an assumption. Nothing is uploaded in bulk, no index of your data is built, and nothing runs when the terminal is closed. Where an organisation chooses to connect other systems (next section), the same loop extends to them under each user's existing permissions. On the plans used for organisational work, the server-side copy of each session is deleted within 30 days and is never used to train models — and organisations with stricter requirements can go further: a **zero-data-retention arrangement**, negotiated directly with Anthropic for qualifying Enterprise accounts, means sessions are not stored at all (detail in the retention section below).

WHAT LEAVES THE MACHINE — AND WHAT NEVER DOES

Transmitted (encrypted, TLS 1.2+; compatible with corporate VPNs and proxies):

- The user's prompts and the model's responses
- The contents of files Claude Code reads for the task in hand
- Operational telemetry — performance metrics that contain **no file content, code or file paths**. All non-essential transmission (telemetry, error reports, feedback commands) can be switched off with a single setting, deployed centrally so individual users configure nothing.

Never transmitted:

- × Files outside the folder the user is working in, unless explicitly granted
- × Files your organisation places under deny rules — credentials, key files and other named paths that Claude Code's tools are forbidden to open
- × Anything from other applications, and nothing to any third party beyond Anthropic — unless your organisation deliberately connects a service (next section)

And on Anthropic's side, your data is never: used to train models (contractual on commercial plans), sold, used for advertising, visible to any other customer, or added to any cross-customer index. There is no mechanism by which one customer's data reaches another.

Connecting other systems (MCP integrations)

Claude Code can optionally connect to other systems — email, cloud drives, calendars, line-of-business tools — through MCP, the open standard for AI-tool connections. Three facts govern the security review:

1. **Default-closed.** Out of the box there are no connections; each one is deliberately added.
2. **Per-user sign-in.** A connector authenticates as the individual user (typically OAuth) and can reach only what that user's account can already reach. The agent inherits existing permissions — it does not create new ones. If a connected tool over-shares, that visibility existed before any AI was attached; the fix belongs in the source system.
3. **Centrally governable.** Managed settings can allowlist approved connectors or block them outright, so the reachable surface is an organisational decision, not a user one.

Content a connector retrieves enters the session and follows the same transmission and retention rules as everything else in this document. The connected service itself continues to process its own data under its own terms, as it did before.

WHERE DOES IT GO, AND FOR HOW LONG?

Processing and storage happen on Anthropic's infrastructure; data at rest is stored in the **US**, encrypted with AES-256. Usage-based Enterprise plans can additionally pin all inference to US-based servers.

Server-side retention of Claude Code session data, by plan:

PLAN	RETENTION	TRAINED ON?	LEGAL BASIS
Personal subscription (Pro/Max)	~30 days — *provided* the user has turned the training setting off (it defaults to on)	**Only if the training toggle is on**	Consumer policy — Anthropic is the data controller
Team	30 days, then deleted	**Never** — contractually prohibited	Commercial Terms + DPA — Anthropic is your data processor
Enterprise	30 days, then deleted	**Never**	Same
Enterprise with Zero Data Retention	**Nothing stored at rest** — prompts and responses are processed in real time and not retained	**Never**	Same

Three honest footnotes that belong in any accurate review:

- **Three exceptions apply on every plan — including under a zero-data-retention arrangement, where nothing else is stored:** sessions flagged by automated trust-and-safety systems may be retained up to 2 years; feedback a user actively submits to Anthropic is retained 5 years (the feedback commands can be disabled organisation-wide); and legal holds apply as they would to any processor.
- **Zero Data Retention (ZDR) is not a plan feature.** It is a per-organisation arrangement available to qualified Enterprise accounts, enabled by Anthropic's account team after an eligibility review — raise it with Anthropic (or with us) if your data classification requires it. It is also available for API-based deployments. One caveat: the frontier models Anthropic designates as "Covered Models" under its safety framework require a 30-day retention window and sit outside ZDR; other models are unaffected.
- **The personal-subscription row is why organisations should not roll out on personal accounts.** Its protections are settings and policy rather than contract — the consumer terms carry no confidentiality obligations, and policy can change (the training default did, in September 2025). Commercial plans make the same protections contractual. (If individuals explore on personal accounts first, the training toggle must be off.)

WHO IS LEGALLY RESPONSIBLE FOR THE DATA?

On Team and Enterprise plans, **your organisation is the data controller and Anthropic is your data processor**, under a Data Processing Addendum incorporating the EU Standard Contractual Clauses and UK and Swiss addenda. The commercial terms are short and unambiguous on the points that matter:

- **"Anthropic may not train models on Customer Content."** The only exception is the Development Partner Programme, which requires an organisation admin's explicit opt-in.
- **You own your inputs and outputs.** Anthropic assigns its rights in outputs to the customer — an IP assignment, not a licence.
- **Customer content is classified as confidential information** with explicit protection obligations.
- **At end of contract**, customer data is returned or deleted within 30 days.
- Anthropic also provides **uncapped copyright indemnification** on paid commercial plans — a point your legal team will note.

Deeper legal-evaluation material (IP nuances, DPIA inputs, indemnity exclusions) is available on request.

WHAT CAN YOU CONTROL CENTRALLY?

Everything below deploys from the centre; users cannot override it.

Machine-level policy — managed settings. Pushed via MDM/GPO (or server-managed without device management), taking precedence over all user settings: permission rules and file-access deny lists (for example, blocking credential and key files from ever being read), command restrictions, network domain allowlists, forced login to your organisation's account, telemetry disabled, and MCP connectors allowlisted or blocked.

Identity and the leaver problem. SSO with domain capture from the Team plan up; SCIM automated provisioning and deprovisioning on Enterprise. Access ends when the account does, and organisational data belongs to the organisation — the common concern about staff on individual accounts taking data access with them does not arise.

Oversight (Enterprise). Audit logs, usage analytics, spend controls, and a Compliance API giving programmatic access to organisational usage data for monitoring and eDiscovery — the governance layer regulated industries ask for.

Data controls. Custom retention policies for organisation content (Enterprise), the US-only inference option (usage-based Enterprise), and ZDR for qualified accounts. Stated plainly: the standard 30-day operational window is not configurable downwards except via ZDR — there is no setting between 30 days and zero.

On the machine itself. Claude Code is read-only by default, asks approval before acting, is scoped to the folder it is started in, and supports OS-level sandboxing — with credential and key files excludable from its reach entirely via the deny rules above. The final control is the person: every action is proposed and approved, and the training builds that supervision into how people work rather than leaving it to policy.

DEPLOYMENT OPTIONS

Standard SaaS (Team or Enterprise) suits most organisations: the controls above, contractual data protection, fastest path to running.

Cloud marketplace deployment suits organisations whose data must stay inside their existing cloud security perimeter. Claude runs on your cloud provider's managed infrastructure — **Anthropic never sees your inputs or outputs**, the cloud provider is the data processor, and your existing compliance frameworks, identity management and (optionally) private networking apply. Billing goes through the cloud agreement you already have.

CLOUD PLATFORM	SERVICE
AWS	Amazon Bedrock
Google Cloud	Vertex AI
Microsoft Azure	Azure Foundry

Claude is available across all three major cloud platforms, so this route is open whichever one your organisation already uses.

CERTIFICATIONS, AND WHO HAS ALREADY DONE THIS REVIEW

CERTIFICATION	SCOPE
SOC 2 Type II	Security and operational controls, independently audited
ISO 27001:2022	Information security management
ISO/IEC 42001:2023	AI management systems
HIPAA	Business Associate Agreements available
FedRAMP High	US government sensitive data (via AWS GovCloud)

Full reports are available through Anthropic's Trust Centre.

The most data-sensitive sectors have already completed this evaluation and deployed: major banks and wealth managers (Anthropic runs a dedicated financial-services offering; RBC Wealth Management is a named customer), healthcare and pharma (Banner Health, Stanford Healthcare; Novo Nordisk cut clinical study report writing times by ~90% running Claude inside AWS Bedrock), most of the UK's top law firms — a sector whose confidentiality duties exceed almost any corporate standard — and the US federal government, where a GSA agreement makes Claude available across all three branches at FedRAMP High.

THE OTHER TWO TOOLS IN THE WORKFLOW

The training workflow includes two supporting tools. Both are simpler cases than Claude Code.

Obsidian is a local application for viewing and editing the plain-text files Claude Code works with. It has no account, no cloud processing and no telemetry — company data never flows through Obsidian's servers. Its optional sync service is end-to-end encrypted (AES-256). Independently security-audited (Cure53 penetration test, report public). It holds no SOC 2 — a small team with, structurally, nothing server-side to certify. For cautious organisations, Restricted Mode disables all community plugins.

Wispr Flow provides voice dictation and processes audio in the cloud. The evaluation hinges on one setting: **Privacy Mode**, which implements zero data retention — audio is processed and immediately discarded, with nothing stored or used for training. Administrators can enforce it organisation-wide, and for work use it should be treated as mandatory. Wispr Flow holds SOC 2 Type II, offers SSO/SAML and BAAs, and its Context Awareness feature (which reads the active application) should be evaluated separately for sensitive workflows.

GETTING STARTED

For most organisations the path is short: choose Team or Enterprise, forward this document and the Trust Centre to your security team — the SOC 2 report and DPA answer most questionnaires — and, where you want the hardening baseline, push the managed settings described above. Nothing else in this document is a precondition. The deeper options — cloud-marketplace deployment, ZDR, custom retention — are there when your data classification calls for them, not before.

On policy: little new policy is usually needed. The human-oversight requirement most AI policies already contain is how this methodology works by design — every output is reviewed and approved by the person who directed it.

KEY RESOURCES

RESOURCE	WHAT IT COVERS
Anthropic Privacy Centre	Data handling and retention, plan by plan
Anthropic Trust Centre	SOC 2 report access, subprocessors, security documentation
Commercial Terms	The no-training clause, IP assignment, confidentiality
Data Processing Addendum	Processor obligations, SCCs, deletion at end of contract
Claude Code Security	Sandboxing, permissions, what leaves the machine
Claude Code Data Usage	Retention detail and telemetry controls

> The technology is enterprise-ready and the most regulated industries have already moved. **The remaining work is people**: the skills, habits and supervision discipline to use these tools well. Multiply Academy provides hands-on AI productivity training for knowledge workers — multiply.academy.

Multiply Academy® is a registered trademark of Multiply Academy Limited. This document is provided for informational purposes; organisations should consult their own legal and compliance teams on their specific regulatory obligations. Facts verified against Anthropic's published policies, July 2026.