ENTERPRISE GUIDE

Enterprise Data Security & Privacy

A guide for organisations adopting AI productivity tools. How enterprise data is protected and how the most data-sensitive industries are already using AI safely.

February 2026

multiply.academy

Your Data Is Protected by Design

The first question most organisations ask when evaluating AI tools is: what happens to our data?

The answer is straightforward. On a commercial AI plan, the provider is contractually prohibited from training on your data. Your files stay on your computer. Conversations are encrypted in transit and at rest. All non-essential telemetry can be disabled.

This is now the industry standard - Anthropic, OpenAI, and Google all enforce these protections on their commercial tiers. This guide focuses specifically on Anthropic's Claude and Claude Code, the tools used in our training, and walks through exactly how your data is protected at every level.

Consumer vs Commercial: The Distinction That Matters

The single most important decision an organisation makes about AI tools is which tier it operates on. The difference between consumer and commercial AI plans is not a matter of features - it is a matter of legal protection.

In February 2026, a US federal judge ruled that documents generated using a consumer version of an AI tool are neither legally privileged nor protected as work product. The reasoning: an AI tool has no law licence, owes no duty of loyalty, and is not bound by confidentiality obligations. On consumer plans, there are no contractual guarantees about how your information is handled.

On a commercial plan, the picture is entirely different.

Anthropic's Data Handling by Tier

	Consumer (Free, Pro, Max)	Commercial (Team, Enterprise, API)
Training on your data	User's choice (opt-in available - verify it is off in Settings)	Contractually prohibited. No exceptions without explicit opt-in to Development Partner Programme.
Data retention	30 days (no training) / 5 years (training enabled)	30 days standard. Zero Data Retention available (nothing persisted).
Contractual protection	Consumer terms of service	Commercial terms with confidentiality obligations. Customer owns all inputs and outputs.
GDPR compliance	Standard privacy policy	Data Processing Addendum with Standard Contractual Clauses.

> **The bottom line:** On any commercial plan, your data is not used for training. Retention is minimal and can be reduced to zero. The protections are contractual - meaning they are legally enforceable, not just settings-based.

The Tools in This Workflow

The training workflow uses three tools. Their security profiles are very different in weight.

Claude Code (Anthropic) is the AI tool - it processes your work content through Anthropic's servers. This is where the substantive security questions live, and the focus of this document.

Obsidian is a local application for viewing and editing files. Your files are never transmitted to Obsidian's servers - the application simply reads files from a folder on your computer. Optional cross-device sync (Obsidian Sync) is end-to-end encrypted.

Wisprflow (or alternative voice input tools) transcribes speech to text. With Privacy Mode enabled, no dictation data is stored or used for training - zero data retention. Wisprflow holds SOC 2 Type II and ISO 27001 certifications.

What Happens to Your Data When You Use Claude Code

Understanding the actual data flow removes much of the uncertainty. Here is what happens when someone in your organisation uses Claude Code:

What stays on your machine

- **All your files remain local.** Your documents and knowledge base live on your computer. The AI reads files when needed but does not upload your entire folder.
- **No permanent copy is made.** On commercial plans with standard retention, conversation data is held for up to 30 days for operational purposes, then deleted. With Zero Data Retention, nothing is persisted at all.

What is transmitted

- **Prompts and context** are sent encrypted (TLS 1.2+) to Anthropic's servers for processing.
- **Operational telemetry** (performance metrics) may be transmitted, but this contains no content from your files. All non-essential telemetry can be disabled with a single setting.

- **Responses** are returned encrypted and rendered locally.

What is not transmitted

- Files outside the folder you're working in (unless explicitly granted access)
- Passwords, credentials, or system configuration
- Data to any third party beyond Anthropic

Encryption

All data is encrypted both in transit (TLS 1.2+) and at rest (AES-256). This is the same standard used by online banking, cloud storage, and enterprise productivity suites.

Enterprise Security: Layers of Protection

Enterprise AI security is not a single feature - it is a stack of complementary protections. Each layer addresses a different risk.

Layer 1: Contractual

Anthropic's commercial terms explicitly prohibit training on customer data. Customers own all inputs and outputs. Customer content is classified as confidential information with explicit protection obligations.

Data Processing Addendums (DPAs) with Standard Contractual Clauses address GDPR requirements. These are incorporated automatically into commercial terms of service.

Layer 2: Certifications

Independent third-party audits verify that security controls work as claimed.

Certification	What It Covers
SOC 2 Type II	Security and operational controls (audited)
ISO 27001:2022	Information security management
ISO/IEC 42001:2023	AI management systems
HIPAA	Business Associate Agreements available on request
FedRAMP High	US government sensitive data (via AWS GovCloud)

The full SOC 2 Type II report is available through Anthropic's Trust Centre on request.

Layer 3: Technical Controls

Anthropic's Enterprise plan includes a full suite of administrative controls:

- **SSO/SAML 2.0** - Single sign-on with your existing identity provider
- **SCIM** - Automated user provisioning and deprovisioning
- **Role-based access control** - Admins assign permission levels
- **Audit logs** - Trace all user actions, system events, and data access
- **Compliance API** - Real-time programmatic access to usage data for regulated industries
- **Custom data retention controls** - Configurable retention and selective deletion
- **Managed policy settings** - Enforce permission standards across your entire organisation

Anthropic also offers a Team plan with SSO and commercial data protection terms, suited to smaller groups. The full administrative controls above - SCIM, audit logs, Compliance API - require the Enterprise plan.

Layer 4: Architectural

Claude Code adds tool-specific protections on top of the platform security:

- **Read-only by default** - The AI cannot modify your files without explicit approval
- **Folder-scoped access** - The tool can only access the folder where it was started
- **Sandboxing** - The AI operates within a secure boundary that prevents it from accessing files or networks outside its permitted scope
- **Restricted operations** - Potentially risky operations are blocked by default
- **Prompt injection defences** - Built-in protections prevent external content from manipulating the AI's behaviour

Deployment Options: Match Your Security Posture

Organisations have a spectrum of deployment options. Each step keeps more within your control. The right choice depends on your data sensitivity, regulatory requirements, and existing infrastructure.

Standard SaaS (Enterprise Tier)

Best for: Most organisations. Professional services, corporate teams, non-regulated industries.

Your team uses Claude directly through Anthropic's platform with enterprise controls. On the Enterprise plan, this includes SSO, audit logs, compliance monitoring, and managed policies across your organisation. Data is encrypted in transit and at rest. Training on your data is contractually prohibited.

Cloud Marketplace Deployment

Best for: Organisations wanting data to stay within their existing cloud security perimeter.

Anthropic has deployed the Claude model onto AWS, Google Cloud, and Azure infrastructure. When your organisation uses Claude through one of these platforms, the processing happens on the cloud provider's managed infrastructure - not on Anthropic's servers. Anthropic never sees your inputs or outputs. You are billed through your existing cloud provider, and their compliance frameworks, networking controls, and identity management apply.

There is no infrastructure to provision or manage. Your IT team enables access through the same cloud console they already use.

Anthropic is the only frontier AI provider available across all three major cloud platforms:

Cloud Platform	Service
AWS	Amazon Bedrock
Google Cloud	Vertex AI
Microsoft Azure	Azure Foundry

For organisations with stricter requirements, cloud marketplace deployment can be combined with private networking - ensuring that traffic between your systems and the AI model stays entirely within your private network and never touches the public internet. Combined with Zero Data Retention, this provides near-air-gapped security while using fully managed cloud infrastructure.

How Regulated Industries Are Using Claude Today

The most data-sensitive industries in the world are already using Claude - through enterprise channels and with appropriate safeguards. These organisations use various Claude products (the web platform, the API, Claude Code, and industry-specific offerings), but the security framework is identical across all of them: the same commercial terms, the same certifications, and the same data handling described in this document.

These organisations worked through the same security and compliance questions and concluded that Anthropic's enterprise protections meet their requirements.

Financial Services

A 2024 Bank of England/FCA survey showed 75% of financial services firms already using AI. Major banks initially banned consumer AI tools over data privacy concerns, then moved to enterprise-controlled adoption.

Anthropic launched Claude for Financial Services in mid-2025, integrating with data providers including FactSet, Snowflake, PitchBook, and S&P Global. Anthropic's partnership with Salesforce delivers Claude to regulated financial services organisations, with RBC Wealth Management among the named customers. Accenture's partnership with Anthropic is jointly developing Claude-based offerings specifically for financial services compliance workflows.

The FCA has confirmed it will not introduce AI-specific rules, instead relying on existing frameworks including the Consumer Duty and Senior Managers and Certification Regime.

Healthcare

Anthropic launched Claude for Healthcare in January 2026, with HIPAA-ready infrastructure and native integrations to medical databases including CMS Coverage Database, ICD-10 codes, and PubMed.

Named healthcare organisations using Claude include Banner Health, Stanford Healthcare, Novo Nordisk, Sanofi, AbbVie, and Genmab. The use cases span clinical documentation, regulatory submissions, and clinical trial analysis.

Novo Nordisk's results are particularly striking: they built a regulatory documentation platform using Claude on AWS Bedrock that reduced clinical study report writing times by 90%. Work that previously required more than 50 people over several months is now handled by three people with AI support. The resulting documents consistently receive positive feedback from regulators.

46% of US healthcare organisations are currently implementing generative AI. The NHS 10 Year Health Plan (July 2025) identifies AI as one of five transformative technologies.

Law Firms

75% of the UK's top 20 law firms now actively promote their AI capabilities to clients. The SRA authorised the UK's first AI-driven law firm in 2025, subject to strict conditions on client confidentiality and hallucination prevention.

Law firms operate under some of the strictest confidentiality obligations of any profession. The fact that the legal sector is among the heaviest adopters of enterprise AI demonstrates that commercial data protection terms - contractual training prohibition, encryption, audit controls - satisfy even legal-grade confidentiality requirements.

Government

Anthropic's Claude for Government is certified for FedRAMP High - the most stringent level for handling unclassified sensitive government data. In August 2025, the US GSA announced a deal making Claude available to all three branches of federal government. The UK Ministry of Defence is working on cloud hosting at Secret and Above Secret classification levels.

The Regulatory Landscape

The regulatory environment for enterprise AI is evolving rapidly across the UK, EU, and US. Existing data protection frameworks (UK GDPR, EU GDPR, HIPAA) already apply to AI tools, and AI-specific regulation is being layered on top - most significantly the EU AI Act, which reaches full implementation in August 2026.

Rather than prescribing specific technologies, regulators are converging on a common set of expectations:

1. **Transparency** - Document what AI tools are used and how
2. **Audit trails** - Maintain logs of AI interactions for compliance monitoring
3. **Human oversight** - Ensure humans review and approve AI outputs
4. **Data protection** - Use tools with appropriate contractual and technical safeguards
5. **Staff training** - Build AI literacy across the organisation

Requirement 5 is now a legal obligation. The EU AI Act's Article 4 (AI Literacy), already in effect since February 2025, requires organisations to ensure their staff have sufficient AI literacy. This applies to any organisation operating in or serving EU markets.

The tools and methodology described in this document address these requirements at different levels. The Enterprise plan provides audit trails (2) and data protection (4) through the platform features described above.

Requirements 3 and 5 are people challenges, not technology challenges - and this is what Multiply Academy delivers. Our training methodology builds human oversight (3) into every task: the core discipline is that every AI output is reviewed and approved by the person who created it before it is used. The curriculum directly addresses staff training (5), giving teams the AI literacy and practical skills to use these tools effectively and responsibly.

Getting Started: Enterprise AI Adoption Checklist

For organisations evaluating AI productivity tools, here are the practical steps:

1. Choose the right tier

- Individual subscriptions (Pro, Max) are fine for initial exploration. Ensure the "Help improve Claude" setting is turned off (Settings at claude.ai) so your data is not used for training.
- Before rolling out across an organisation, move to a commercial plan (Team or Enterprise) for contractual data protection - legally enforceable, not just a setting.

2. Engage your IT and security teams

- Share Anthropic's Trust Centre documentation and SOC 2 report.
- Review the DPA and confirm it meets your data protection requirements.
- The three-sentence summary they need: "On any commercial plan, Anthropic is contractually prohibited from training on your data. They hold SOC 2 Type II, ISO 27001, and ISO 42001 certifications. All non-essential telemetry can be disabled."

3. Match deployment to data sensitivity

- Standard enterprise tier: suitable for most organisations
- Cloud marketplace (Bedrock, Vertex, Azure Foundry): for data that must stay within your cloud perimeter, with optional private networking for regulated industries requiring zero public internet exposure

4. Configure enterprise controls (Enterprise plan)

- Enable SSO/SAML for single sign-on
- Set up SCIM for automated user provisioning
- Configure audit logging and Compliance API for monitoring
- Deploy managed policy settings across all users
- Disable non-essential telemetry if required (a single managed setting deployed centrally - individual users do not need to configure anything)

5. Establish governance

- Document your AI tool usage policy
- Define which information categories are in scope for AI processing and which are excluded (e.g. passwords, financial credentials, specifically regulated data)
- Train staff on responsible AI use. Multiply Academy provides structured training that covers both AI literacy and the practical supervision skills needed to use these tools effectively and safely

Key Resources

Resource	What It Covers
Anthropic Privacy Centre	Data handling, retention, certifications
Anthropic Trust Centre	SOC 2 report access, security documentation
Claude Code Security Docs	What data leaves the machine, sandboxing, telemetry

Claude Code Data Usage	Detailed data handling for Claude Code
Anthropic Commercial Terms	Contractual data protection commitments

Next Steps

The technology is enterprise-ready. The security protections are in place. The regulated industries have already moved. The remaining challenge for most organisations is the people side: building the skills, habits, and governance to adopt AI effectively.

Multiply Academy provides hands-on AI productivity training for knowledge workers - from first principles through to full workflow adoption. Learn more at multiply.academy.