SHIFT AI
SOLUTIONS

17 August 2026

AI Data Security & Privacy

Client Assurance Document & Service Guarantee

Prepared by Shift AI Solutions

This document explains exactly how your data is protected when AI is deployed into your business by Shift AI Solutions — and provides our written commitment to your security.

30 Days	0	100%
Maximum data retention on Anthropic API servers	Times your data is used to train AI models	Commercial-grade terms on every deployment

1. The Critical Distinction: Consumer vs. Business AI

Not all AI accounts are equal. There are two fundamentally different tiers, and they treat your data in completely different ways. This is the single most important thing to understand about AI security in a business context.

Feature	Personal AI accounts The free or personal-subscription apps your staff might sign up for themselves	✓ Business AI accounts Commercial accounts, contracted to your business — what we set up for you
Is your data used to train the AI?	YES — by default	NEVER
How long is your data kept?	Up to 5 years	30 days maximum
Who owns and controls it?	Anthropic	Your business (you are Controller)
Meets the Australian Privacy Act?	Only if opted out	Yes — by design
Safe for health records?	No	Yes (with BAA agreement)
Can you see who used it, and when?	No	Yes — configurable
Safe for business use?	Risky	Yes — built for it

Key takeaway: When Shift AI Solutions deploys Claude into your business, we exclusively use commercial-tier API or Enterprise accounts. Your data is never used to train any AI model, and it is automatically deleted from Anthropic's servers within 30 days.

Zero-retention option: For clients with stricter requirements, Anthropic also offers a zero-data-retention arrangement, where prompts and responses are not stored at all after the response is returned. We can arrange this where your compliance obligations call for it.

2. What Actually Happens to Your Data

Understanding the data flow helps you explain the security posture to your own stakeholders, IT team, or compliance auditors.

Data Flow — Cloud API Deployment (Standard)

1	Employee enters a prompt or query into the AI tool.
2	Data travels to Anthropic's servers over encrypted HTTPS — the same encryption standard used by online banking.
3	Claude processes the request and returns a response.
4	The interaction is logged temporarily for operational purposes only.
5	After 30 days, all inputs and outputs are permanently deleted from Anthropic's servers — automatically, with no action required from you.
6	At no point is your data used to train, fine-tune, or improve any AI model. This is contractually guaranteed under Anthropic's Commercial Terms.

Anthropic's Role: Processor, Not Controller

Under Anthropic's Commercial Terms of Service, your business is legally defined as the Data Controller — meaning you own and control your data. Anthropic acts only as a Data Processor, handling your data strictly to deliver the service, nothing more. This is a critical legal distinction under the Australian Privacy Act, GDPR, and other regulatory frameworks.

On-Premise / Local Deployment Option

For clients with the highest security requirements — such as those handling classified information, sensitive financial records, or health data — we can architect a local deployment where all AI processing happens entirely within your own infrastructure. In this configuration, zero data ever leaves your network. This meets GDPR, HIPAA, and Australian Privacy Principles by design.

3. Honest Risk Disclosure & Our Mitigations

We believe in full transparency. Any technology carries some level of risk, and AI is no different. Below we outline the known risk categories and the specific controls Shift AI Solutions puts in place to address each one.

Prompt Injection

Threat: A malicious actor crafts inputs designed to trick the AI into revealing information or behaving unexpectedly.

Our Control: We implement input validation, sandboxing, and output filtering as part of every deployment. System-level instructions prevent the AI from discussing out-of-scope topics or revealing system configuration.

Sensitive Data in Prompts

Threat: Employees may inadvertently paste passwords, personal identifiers, or confidential client data into AI queries.

Our Control: We deliver staff training as part of onboarding, establish clear Acceptable Use Policies, and can implement automated PII detection that strips sensitive fields before they reach the AI layer.

Access Control Gaps

Threat: Without proper permissions, all employees may have access to AI tools connected to sensitive business systems.

Our Control: We design role-based access controls from day one — ensuring the right people have access to the right data, with audit trails for every interaction.

Data Leakage Between Users

Threat: In shared deployments, one user's queries could theoretically surface another user's information.

Our Control: Our deployments use isolated conversation contexts and session management to ensure complete separation between users and departments.

Regulatory Non-Compliance

Threat: Using consumer-grade AI for business purposes may violate Australian Privacy Act obligations, GDPR requirements, or industry-specific regulations.

Our Control: We deploy exclusively under commercial terms, provide a Data Processing Addendum on request, and can assist with compliance documentation for regulatory auditors.

4. Regulatory Compliance Framework

Shift AI Solutions deployments are designed to operate within the following regulatory frameworks applicable to Australian businesses:

Australian Privacy Act 1988 & APPs	Our deployments operate under commercial API terms, where your business remains the Data Controller. Data is not retained beyond operational necessity (30-day maximum), supporting compliance with the Australian Privacy Principles including data minimisation and purpose limitation.
GDPR (EU — for businesses with EU clients)	Commercial terms include a Data Processing Addendum (DPA) available on request. Data subject rights (access, erasure, portability) are supported. Anthropic acts as a compliant Processor under Article 28.
HIPAA (Healthcare — where applicable)	For healthcare clients, we configure deployments under Anthropic's HIPAA-eligible settings, including a Business Associate Agreement (BAA). Web search features are disabled under HIPAA configuration for full compliance.
PCI-DSS (Financial / Payment Data)	Card data and financial identifiers are explicitly excluded from AI prompt inputs via our data handling policies and, where required, automated filtering to prevent PCI-DSS scope creep.
ISO 27001 Alignment	Our deployment methodology aligns with ISO 27001 information security principles including access control, audit logging, incident response planning, and vendor risk management.

5. Our Secure Deployment Stack

Every Shift AI Solutions deployment includes the following security controls as standard. These are not optional add-ons — they are part of how we build.

01 Commercial API / Enterprise Account

We exclusively use Anthropic's commercial-tier access for all client deployments. This is non-negotiable and ensures your data is never used for model training.

02 Role-Based Access Controls (RBAC)

Users are granted access only to the AI tools and data sources relevant to their role. Department-level isolation prevents cross-contamination.

03 Encrypted Data Transit

All communication between your systems and the AI layer uses TLS 1.2+ encryption — the same standard as online banking.

04 Audit Logging

Every AI interaction is logged with user ID, timestamp, and query category, for example “invoice query” or “quote drafting”. We log the type of request, not the words your staff type, so you get accountability without monitoring your team's messages. Logs are retained per your internal IT policy and available for compliance review.

05 Acceptable Use Policy (AUP)

We provide a customised AUP for your business and deliver staff training to ensure employees understand what data should and should not be shared with AI systems.

06 Incident Response Plan

In the event of a suspected data issue, we provide a documented escalation and response process aligned with the Australian Notifiable Data Breaches scheme.

07 Regular Security Reviews

We conduct quarterly reviews of your AI deployment configuration to ensure it remains aligned with evolving security standards and your business needs.

6. Your Data During the Assessment

Before any AI is deployed, the assessment itself involves us handling your people's words. Here's exactly how that works.

What we collect

Audio recordings of each interview, written transcripts produced from them, and our own notes. Nothing else — we do not access your systems, files, or customer data during the assessment unless you specifically ask us to.

Consent comes first

We confirm every participant is comfortable being recorded at the start of their session. Anyone who would rather not be recorded simply says so, and we take written notes instead. No explanation needed, and it makes no difference to how we treat their input.

Who hears the recordings

Only the Shift AI Solutions consultants running your assessment. Recordings and individual transcripts are never shared with your management team. This is deliberate: people describe how work actually happens, including the workarounds, only when they know their manager is not reading the transcript. That candour is what makes the assessment worth doing.

Where it is stored

Recordings are transcribed on our own equipment using a local speech-to-text model. The audio is never uploaded to a transcription service, never sent to a cloud provider, and never leaves our device. Transcripts are then processed through Anthropic's commercial API under the terms described earlier in this document — never a consumer account, and never used to train any model.

How it appears in your report

Findings are evidenced with direct quotes, because that is what makes them credible. Quotes are attributed by role, not by name (for example, "a senior driver told us..."). Where a quote could only have come from one person, we either paraphrase it or check with that person first.

How long we keep it

Audio recordings are kept only until the transcript has been checked for accuracy, then permanently deleted. Transcripts and analysis are retained for 6 months after the assessment so we can support you through implementation, then permanently deleted. You can ask us to delete everything sooner at any point, and we will confirm in writing once it is done.

7. The Shift AI Solutions Security Guarantee

OUR WRITTEN COMMITMENT TO YOU

Shift AI Solutions guarantees the following for every client engagement:

✓ We will ONLY deploy AI systems using commercial-grade API or Enterprise accounts — never consumer accounts — ensuring your data is never used to train AI models.

✓ Your data will be automatically deleted from AI provider servers within 30 days of each interaction. No long-term storage, no data hoarding.

✓ We will implement role-based access controls, audit logging, and staff training as standard components of every deployment — not optional extras.

✓ We will provide you with full transparency about what data flows where, who can access it, and how it is protected — in plain language, at any time you ask.

✓ If a security concern arises, we will respond within 24 hours and provide a written incident report within 72 hours, in accordance with the Australian Notifiable Data Breaches scheme.

✓ We will never recommend an AI solution that we would not be comfortable using ourselves with our own most sensitive business data.

This guarantee forms part of our client service agreement and is provided in writing as a condition of every Shift AI Solutions engagement.

Shift AI Solutions	Document Date	Version
shiftaisolutions.net	17 August 2026	1.0 — Current

Questions about this document? Contact us at shiftaisolutions.net