PROMPT WORKS STUDIO

Chaos to calm, built to be yours.

Campbell Street Home Care

AI use policy

Prepared for the Director

Chad Maxwell · Prompt Works Studio · Toowoomba

November 2026

> **A made-up example.** Campbell Street Home Care and its people aren't real. The software, its terms and the rules in it are real, as read on the dates shown.

PROMPT WORKS STUDIO

Part 1: How this policy was built

This is your final policy, after our call on Friday 6 November. Part 1 records what I found, what you decided and why, so when the Director reviews it next November, the reasons behind each rule are here. Your business had a Workflow Audit with me first, so the policy draws on that too. Part 2 is the policy your team signs.

It was written with AI (Claude) and checked by me. It isn't legal advice.

What went into it

Step	When	What it gave the policy
Workflow Audit, a day and a half in your office	Tuesday 27 and Wednesday 28 October	The five office jobs that take the most time, the AI already in use, and the fix-now list
Staff survey, anonymous, sent to all 42 staff, open five working days	Wednesday 28 October to Tuesday 3 November	Which tools your team uses, what goes into them, and what they'd like help with
Interview with the Director and the Intake and admin officer, about an hour	Wednesday 4 November	How each job runs, who checks what, and your choices on the hard questions
Each tool's terms read on the maker's own pages	Before the call, and again on Sunday 8 November for this copy	Which tools can take client details (the appendix)
Draft, then our 30-minute call	Thursday 5 November, Friday 6 November	Your decisions below, and one round of changes

What I found, and what the policy does about it

What I found	What the policy does	Section
Incident write-ups and family emails were being tidied in a free chatbot on a personal phone, with clients' names and health details in them	Personal accounts are out, and incident records stay out of AI	5, 6, 7, 8
ShiftCare's own AI was in use, with no list of what's on	A register. Only the ShiftCare note tools are on	4
Nobody was named for AI, and nobody was shown how to use it safely	An AI lead, and everyone is shown the tools and signs	3, 10
Consent forms and the privacy policy don't mention AI or information going overseas	New consent wording and privacy policy lines. No client's details go into AI until their file says "AI OK"	7, 11, 15
Care plan reviews take about 7 office hours a week, and review dates sat in a spreadsheet	The dates were moved into ShiftCare at set-up. Claude builds review packs and drafts plans and letters, and a care partner reads every line	4, 5
The ShiftCare incident form was missing witnesses' contact details, whether the incident's reportable, and who was consulted	The form now asks for all 13 details. Deciding and lodging stay with the Director	8, 17

What I found	What the policy does	Section
From 10 December 2026 the OAIC's automated-decisions duty starts, and it lists aged care	Claude helps draft care plan updates, so the privacy policy carries the line now, worded by your adviser	11, 15
Busy times: month end and the weeks before each quarterly claim	No new tool or use is approved then	3

What your team said

30 of your 42 staff answered the survey: all 8 in the office and 22 of the 34 care workers. With six or more, you saw the answers as written.

They add up like this. About a third of the care workers who answered use ShiftCare's voice-to-text or note rephrase, and two said they'd used a free chatbot on their own phone for work. In the office, most want help with care plan reviews and letters. Most asked what they're allowed to put in. That's why section 7 lists it, and why every job in section 5 says who checks.

Your decisions

You made these on our call. Where you went a different way from my recommendation, the reason's recorded.

Question	My recommendation, and why	What you decided
What never goes into AI	Incident records, SIRS notices and complaints, ID documents, Medicare, My Aged Care, Centrelink, bank and card numbers, logins, photos or recordings of a client, staff pay and personal details, and anything a client asks to keep out.	As recommended
What AI is never used for	Deciding if an incident's reportable, what's in a care plan, clinical care, the figures on a statement or claim, and who gets a service	As recommended
Client details in Claude	Yes, on the Team plan, for care plan reviews only, for clients marked "AI OK". Anthropic won't train on business plan chats by default	As recommended
Incident records and notices	Keep AI out. The Rules put the decision on you	As recommended
ShiftCare's SmartMatch and Vacant Visit Agent	Leave off. They choose who visits whom, which could be a decision about a person	As recommended
ShiftCare's Care Signals and Visit Summaries	Leave off. ShiftCare's AI page lists Care Signal Detection as in development, and its pricing page doesn't show Visit Summaries in a plan	As recommended
Microsoft 365 Copilot in Word and Outlook	Not approved. Claude already covers care plan reviews, and no other job needs client details in a second AI tool	As recommended. Copilot Chat on your business sign-in is approved for general wording with no client details
Where AI work is recorded	A line in the client's ShiftCare file: the tool, what it did and who checked it	As recommended

Question	My recommendation, and why	What you decided
Existing clients' consent	Ask each client at their next care plan review, so it's spread across a year	**Changed:** every client is asked by phone and a new consent form by Friday 18 December 2026, because the privacy policy changes now and you want everyone settled together
Who's the AI lead	The Director, who's accountable	**Changed:** the Intake and admin officer keeps the register day to day, since the Director is on call as clinical lead and the Intake and admin officer already sets people up in ShiftCare. The Director stays accountable
Training for care workers	A 30-minute session for each of the 34 care workers	**Changed:** a 15-minute talk at each team meeting plus a one-page card, since pulling workers off visits into one room costs more in cover
Breaking the policy, and owning up	Handled like any other breach of your rules, and owning up isn't punished, so people tell you early	As recommended

Rules for your field

- **Aged Care Act 2024**, in force since 1 November 2025. Section 168: use a client's information only for their care or the reason it was given, share it only with consent or for their care, and keep reasonable safeguards. Sections 7, 11 and 17.
- **Aged Care Rules 2025.** Information about a person is accurate, current and confidential in line with their informed consent (Outcome 2.7). Incident records need 13 set details, and notices have deadlines. Sections 5, 8 and 17.
- **The Privacy Act.** A provider that gives a health service and holds health information is covered whatever its turnover, so you're covered. From 10 December 2026 your privacy policy must say if a computer program makes, or substantially and directly helps make, a decision that could significantly affect someone. A person checking doesn't on its own take a decision out of it, and the OAIC says to be cautious, so your privacy policy says now that Claude helps with care plan reviews (section 15).
- **The OAIC** recommends keeping personal information, and particularly sensitive information, out of public AI tools, and telling people in your privacy policy how you use AI. Sections 6, 7 and 11.
- **The Aged Care Quality and Safety Commission.** I found no AI guidance for providers on 8 November. My emails will say if that changes.
- **The Fair Work Commission.** You have staff, so section 14 covers its rule on AI in case documents, in force since 20 October 2026.

What's not in it, and why

- **AI agents** that send, book or choose on their own. ShiftCare has some, and they're off.
- **AI note-takers** that record visits, calls or meetings. Nobody does, so they aren't approved.
- **AI in clinical care**, such as falls prediction. You don't use any.
- **Copyright** in AI-made work. You don't use AI for marketing or creative work.

If any of that changes, my emails at 3, 6 and 12 months cover it. I can also make changes at a yearly review for $350 inc GST.

Part 2: The policy

AI use policy, Campbell Street Home Care

Policy owner	Director
AI lead	Intake and admin officer
Version	1.0
Adopted	Friday 6 November 2026
Next review	November 2027
Current copy	Office: the Policies folder in our Microsoft 365. Care workers: a printed copy from the office, and the one-page card. A printed copy may be out of date, so check the version first
Goes with	Our privacy policy, our consent form and service agreement, our medication policy and our incident management system

This policy was drafted with AI (Claude) by Prompt Works Studio, then checked and approved by Chad Maxwell. Campbell Street Home Care has read it, changed it to fit and adopted it as its own. It isn't legal advice.

1. Why we have this policy

About 120 people in Toowoomba, Highfields, Oakey and Pittsworth trust us with their care and their health details. We use AI to take some paperwork off our care partners, so they've got more time with clients. This policy sets out how, so our records stay right and each person's information stays private.

2. Who it covers, and what counts as AI

It covers everyone who works for us or with us, office team and care workers, casual or permanent, on any device used for our work, including your own phone.

It also covers anyone we bring in to help, like our IT provider. They don't use AI on our clients' information unless it's a tool on our register, used the way the register allows, and the Director has said yes in writing.

AI means any tool that writes, sums up, suggests or decides things for us. That includes chatbots like Claude and ChatGPT, and the AI inside software we already use: ShiftCare's voice-to-text, note rephrase, Completeness Coach, Visit Summaries, SmartMatch, Vacant Visit Agent, Care Signals and the rest, Copilot Chat and Copilot in Word and Outlook, and Xero's JAX and suggested coding. It doesn't include ordinary spreadsheet formulas, fixed rules like a required field on a ShiftCare form, or plain reports.

3. Who's in charge

- **The Intake and admin officer is our AI lead.** They keep the register in section 4, show staff how to use our tools, and are who you tell when something goes wrong.
- **The Director is accountable** for all work that leaves the business, however it was made, and signs off this policy.
- **Each month the Director looks at three client files** where Claude helped with a review, to check the AI work was recorded and checked as this policy says.
- **No new tool or new use goes on the register at the end of a month or in the weeks before a quarterly claim.** Those are our busy times. A request then waits for the next quiet fortnight.

4. Our AI register

We only use the tools on this list, on the business's own accounts, for the jobs listed. A tool goes on the list only after the AI lead has read what its maker says about training on our data and the Director has agreed. Want a new tool or use? Ask the AI lead.

The register can change between reviews. Every change goes on the review record, and the AI lead tells the team first.

Tool	Account	Approved for	Client details?	Approved
ShiftCare's note tools: voice-to-text, AI note rephrase and Completeness Coach	Our ShiftCare Premium account, each worker on their own login	Care workers' and care partners' visit notes, and a care partner dictating notes after a home visit	Yes, inside ShiftCare only, for clients marked "AI OK" (section 11)	6 November 2026, by the Director
Claude	The Team plan, five seats: the Director, the three care partners and the Intake and admin officer. Run in chat by a named person each time. Nothing runs on a timer	Care plan reviews only: the reviews due from ShiftCare (the dates were moved there at set-up), a review pack from the notes since the last plan, the updated plan wording and the covering letter. Reads ShiftCare through its read-only Claude connection	Yes, for clients marked "AI OK". Never what section 7 rules out	6 November 2026, by the Director
Microsoft 365 Copilot Chat	Our business Microsoft 365 sign-in	General wording, checklists and templates with no client, family or worker details in them	No	6 November 2026, by the Director

Not approved:

- ShiftCare's SmartMatch and Vacant Visit Agent. They choose who visits whom, which could be a decision about a person.
- ShiftCare's Visit Summaries. ShiftCare's pricing page doesn't show them in a plan, so they're not confirmed as part of ours.
- ShiftCare's Care Signals, AI reminder messages, Inbox Signals and AI dashboard summaries. No job needs them, and ShiftCare lists some as still in development.
- Microsoft 365 Copilot in Word and Outlook. Claude covers care plan reviews, and no other job needs client details in a second AI tool.
- Xero's JAX and suggested coding. Statements and claims stay in ShiftCare with the Finance and claims officer, so no job needs them (appendix).
- Any personal or free AI account, including the free ChatGPT. Its terms let it train on what goes in (appendix).
- AI note-takers, or anything that records or transcribes visits, calls or meetings.
- AI agents that send, book, pay, roster or change things without a person each time.
- Any chatbot or AI voice that talks to clients or families for us.
- Any other tool that would make, or help make, a decision about a person. Our privacy policy has to say so first (section 15).

What each tool's terms say about training is in the appendix.

5. How AI is used in our work

Every job has a person who checks the AI's work before it's used. Some steps stay with people for good (section 8).

Job	Where AI helps	Who checks, and against what
Getting a new client started	A care partner dictates the visit notes into the ShiftCare app and ShiftCare's voice-to-text sets them out. The checklist and plain-language agreement template aren't AI	The care partner reads the transcribed notes before using them. The Intake and admin officer checks details against the person's Notice of Decision
Incident records and notices	None. The worker's account goes into ShiftCare in their own words	The care partner checks the record against the worker's account. The Director decides and lodges
Care plan reviews	Claude lists the reviews due, builds a review pack from the notes since the last plan, and drafts the updated plan wording and the letter after the meeting	The care partner reads every flagged note in full and checks every line of the plan against their own notes. The person gets a copy every time it changes
Keeping families in the loop	None. Staff use templates and the Family Portal	Whoever sends it reads it first, and checks what the person has agreed to share
Monthly statements and claims	None	The Finance and claims officer checks every figure against what was delivered
Visit notes by care workers	Voice-to-text, note rephrase and Completeness Coach in ShiftCare	The worker reads it and fixes it so it says what happened. The care partner reads the notes

Right and wrong, in our work

Right. A client's care plan review is due. The care partner checks the file says "AI OK", then asks Claude for a review pack from the notes since the last plan, which they copy in. They read every flagged note in full, talk it through with the person and their supporter at home, then ask Claude to draft the updated plan and letter. They check every line against their own notes, put the final version into ShiftCare, send the person a copy, and add a line to the file: "Review pack and plan drafted in Claude, checked by the care partner."

Wrong. A care worker finds a client on the floor, calls the office, then puts the account into a free chatbot on their own phone "to make it neater", with the client's name. That's three things ruled out: a personal account, health details in it, and an incident record in AI. What to do instead: write the account in ShiftCare in your own words, and tell the AI lead what went into the chatbot. Nobody's in trouble for owning up (section 16).

Wrong. A client's daughter emails asking how her mum's been. A staff member pastes the week's ShiftCare notes into Copilot Chat to draft a reply. Copilot Chat is for wording with no client details in it, and the client may not have agreed to share the notes. What to do instead: check what the person has agreed to share, then answer from a template or point the family to the Family Portal.

6. Personal accounts

Don't use a personal AI account for our work, on your phone or anywhere else, including the free version of a tool we pay for. If you've got client details in one now, tell the AI lead. You won't be in trouble (section 16).

7. What goes in, and what never does

- Our own templates, public information and general questions can go into any tool on the register.
- Client details go in only where the register allows it, only for clients marked "AI OK", and only what the job needs.
- **Never put in any AI tool, Claude included:** incident records, SIRS notices and complaints, ID documents, Medicare, My Aged Care, Centrelink, bank and card numbers, logins and passwords, photos, recordings or video of a client, a worker's pay, health or personal details, or anything a client has asked us to keep out.
- **Clients without "AI OK" on their file:** that covers clients who've said no, and everyone whose new consent form isn't back yet. Their details stay out of every AI tool, ShiftCare's included. Their visit notes are typed by the worker, not dictated or rephrased.
- Nobody sends thumbs up or down feedback on Claude work that has client details in it, because Anthropic may then use that chat.
- If something goes in by mistake, tell the AI lead the same day.

8. What we never use AI for

- **Deciding if an incident is reportable, how serious it is, and lodging the notice.** The Director decides and lodges.
- **What goes in a person's care plan.** AI can set out what the care partner wrote. The care partner decides.
- **The home visit, going through the agreement, and the monthly care management contact.** A care partner does them, with the client's supporter there if they want.
- **Clinical care and medicines.** Medicines are given, watched and reviewed by registered health practitioners, and clinical oversight is the Director's. Care workers' prompts follow our medication policy.
- **Noticing that someone's getting worse.** A summary can't be the only way a change is picked up. Workers tell the care partner or the office straight away.
- **The figures on budgets, statements and claims.**
- **Who gets a service, or which worker is sent.**
- **Anything meant to mislead anyone.**

9. Checking AI work

- The person who checks must know the work well enough to spot a mistake. Section 5 says who.
- Check names, dates, services and prices against the source: ShiftCare, the Notice of Decision, or the worker's own notes. AI can make things up that look right.
- A summary can miss something, so read the flagged notes in full.
- Don't use AI to answer a question you couldn't check yourself. If you're not sure an answer's right, ask.

10. Training

- Before anyone uses a tool on the register, the AI lead shows them how, what it's approved for, and this policy, and they sign the sign-off page. The office team does that on Monday 9 November.
- Care workers get a 15-minute talk at their team meeting and the one-page card, and sign by Friday 20 November. Until they've signed they use no AI tool, ShiftCare's included, and type their notes.
- Every November the whole team goes over this policy again, with any change to the register.
- The AI lead keeps a list of who's been shown what, and when.

11. Telling our clients

- Our consent form and service agreement tell clients and their families how we use AI, say that Claude's data may be stored or processed outside Australia, including in the United States, and that ShiftCare's AI runs in Sydney, and ask for the client's OK. Our adviser worded the clause. It's in every new agreement from Monday 9 November 2026, and every existing client gets it by phone and a new form by Friday 18 December 2026.
- Our privacy policy says we use AI tools, which kinds of information go into them, that personal information is likely to go overseas, and that Claude helps with care plan reviews, which feed decisions about a person's care (section 15). Our adviser worded those lines too.
- When a new consent form comes back agreeing, the Intake and admin officer writes "AI OK" at the top of the client's ShiftCare file. If they say no, or the form isn't back, there's no mark, and their details stay out of every AI tool (section 7).
- The client chooses what their family sees. We share a client's information with a registered supporter where the client has agreed, and with anyone else only with their consent.
- If a client or family asks how AI was used on their care, we tell them.

12. Keeping a record

When AI helps with a client's review, add a line to the client's ShiftCare file: the tool, what it did, and who checked it. It takes a line, and it's how the Director's monthly look at three files works.

13. Fake voices and messages

AI can copy a voice or a face. Before we give out a client's information, change a visit on a caller's say-so, change a client's or a worker's bank details, pay anything new, or give anyone access to ShiftCare, we call back on a number we already have on file, and check the person is the client's registered supporter or has the client's OK. That goes even when it sounds like someone we know well.

14. Fair Work Commission documents

If AI helps write anything for a Fair Work Commission case, the document says AI was used and how, and a person checks it and says they have. Using AI only for spelling, grammar or layout doesn't count. A witness statement is in the witness's own words. Other people's personal details and confidential case information don't go into AI for it at all. The Commission counts paid Claude and Copilot as public AI, so that covers both Claude and Copilot.

15. Privacy and data breaches

- We give a health service and hold health information, so the Privacy Act covers us whatever our size, including the data breach scheme.
- If personal information goes somewhere it shouldn't, through AI or anything else, tell the AI lead the same day. The AI lead and the Director act straight away to contain it.
- The Director then assesses it as quickly as they can, and always within 30 days, to work out whether it's likely to cause someone serious harm. If it is, and we can't stop that harm, we tell the people affected and the OAIC.
- From 10 December 2026 our privacy policy has to say if a computer program makes, or substantially and directly helps make, a decision that could significantly affect someone. Claude drafts care plan updates and review packs, which feed decisions about a person's care, so our privacy policy says so now. It names the kinds of information (care notes and plans, including health information) and the kind of decision (care plan reviews). A person checking doesn't on its own take a decision out of the rule.
- Before we turn on any other tool that would make, or help make, such a decision, the privacy policy gets a line for it first.
- Our privacy policy says personal information is likely to go overseas. Claude's data may be stored or processed outside Australia, including in the United States (Anthropic's privacy centre page, 15 June 2026).

16. When something goes wrong

- Tell the AI lead the same day if client information went somewhere it shouldn't, AI work went out with a mistake in it, or a tool did something unexpected.
- If a mistake in AI work reached a client and could have harmed them, the care partner records it as an incident in ShiftCare, and the Director decides whether it's reportable.
- If ShiftCare or Claude is down, we carry on by hand as we did before.
- Breaking this policy is handled like any other breach of our rules. Owning up to a mistake, or to AI you've already been using, won't get you in trouble. We'd rather know.

17. Our duties as a registered provider

- **Personal information.** We use a client's information only for their care or the reason it was given, share it only with their consent or for their care, and keep reasonable safeguards (Aged Care Act, section 168).
- **Accurate records.** What we hold about a person must be accurate, current, understood by them and their supporters, and confidential in line with their informed consent (Outcome 2.7). That's why a person checks every AI draft against the source.
- **Incidents.** Our ShiftCare incident form asks for all 13 details the Rules list, and we keep each record 7 years. The Director decides whether it's reportable and priority 1 or 2. A priority 1 notice goes to the Commissioner within 24 hours and a priority 2 notice within 30 days. AI has no part in it (section 8).
- **Agreements and plans.** Agreements are in plain language. The person gets a copy every time their plan changes.
- **Changes in a person's condition** are reported up straight away. An AI summary is never the only way we notice (Outcome 3.3).
- **The Commission.** It hasn't published AI guidance for providers that we've found. The Director looks again at each review.

18. Reviewing this policy

The Director reviews this policy each November, in the first two weeks, and sooner if we take on a new tool, something goes wrong, or the law or the Commission's guidance changes, including when the OAIC's duty starts on 10 December 2026. Every review goes on the review record, even when nothing changes. When the policy changes, everyone signs a fresh sign-off page.

Review record

Fill in a row at every review, and whenever the policy or the AI register changes. The AI lead keeps this page with the policy.

Date	What was reviewed or changed	Why	Reviewed by	Team signed again?
Friday 6 November 2026	First adopted	Written after the Workflow Audit and the policy interview	Director	Office team Monday 9 November. Care workers at team meetings by Friday 20 November

PROMPT WORKS STUDIO

Sign-off

Everyone who works for us signs this page: the team in the first weeks after the policy's adopted, and anyone who joins after, before they use any AI tool for our work (section 10). When the policy changes, everyone signs again on a fresh page. The AI lead keeps the signed pages.

Version signed: 1.0, adopted Friday 6 November 2026.

By signing, I confirm I've read this policy, I've been shown the tools I'm approved to use, and I'll follow it.

Role	Name	Signature	Date
Director			
Care partner			
Rostering coordinator			
Intake and admin officer			
Finance and claims officer			
Home care worker			

Appendix: what each tool's terms say about training

Read on each maker's own pages before the register was approved, and again on Sunday 8 November 2026 for this copy. Before approving Claude, the AI lead and the Director also read Anthropic's commercial terms and ShiftCare's AI and data usage policy.

Tool	What the maker says	Where, and when read
Claude, Team plan	"By default, we will not use your inputs or outputs from our commercial products", unless someone sends feedback. The pricing page shows model training as "None by default" on Team and "Opt-out" on Free, Pro and Max. Its "Where are your servers located?" page names Enterprise and the API, not Team, and says "data is stored in the US" and traffic may go to "select countries in the US, Europe, Asia and Australia"	privacy.claude.com, "Is my data used for model training?" (18 August 2026) and "Where are your servers located?" (15 June 2026), and claude.com/pricing, 8 November 2026
ShiftCare's AI (note tools and the Claude connection)	"Your data is never used to train AI models". Its AI and data usage policy adds "without your explicit opt-in", and the AI settings article says it "cannot be enabled", so we haven't opted in. "All AI infrastructure is hosted in AWS Sydney (ap-southeast-2)." For data sent to an AI tool you've connected, "responsibility for the subsequent processing, storage, and protection of that data" is the customer's	shiftcare.com AI data controls and AI and data usage policy (no date shown). help.shiftcare.com, AI settings article, 7 September 2026, and MCP Server Data Handling Overview, 26 June 2026. All read 8 November 2026
ShiftCare's SmartMatch, Vacant Visit Agent, Visit Summaries, Care Signals and other AI (not approved)	The same ShiftCare statements cover them. Not approved because they could decide about a person, no job needs them, ShiftCare's pricing page doesn't show Visit Summaries in a plan, or ShiftCare lists Care Signal Detection as in development	shiftcare.com/ai-and-automation and shiftcare.com/pricing (no dates shown), 8 November 2026
Microsoft 365 Copilot Chat	"Your data isn't used to train foundation models", and "We won't use your data except as you instruct." Web searches go to Bing under separate terms, so no client details go in	learn.microsoft.com, enterprise data protection page, updated 4 November 2026, read 8 November 2026
Microsoft 365 Copilot in Word and Outlook (not approved)	"Prompts, responses, and data accessed through Microsoft Graph aren't used to train foundation LLMs". Not approved because Claude covers the job	learn.microsoft.com, Copilot privacy page, last updated 3 November 2026, read 8 November 2026

Tool	What the maker says	Where, and when read
Xero's JAX (not approved)	Data sent to its AI providers "isn't used to train the LLMs". Not approved because no job needs it. Xero's terms, clause 23, also let it use anonymised data from yours to "develop new services"	xero.com/au, 8 November 2026
ChatGPT, free version (not approved)	"When you use our services for individuals (Free, Plus, Pro), we may use your content to train our models." You can opt out, but each person would have to	help.openai.com, 8 November 2026