Anum

Putting AI into an Australian business, properly and safely, in the right order

A working guide for Australian businesses. Written by Anum, Melbourne.

It opens with five places AI pays first, then sets out the order we use. It starts with the work, not the software. Nine short chapters, one policy you can adapt tonight, and a scoring sheet for picking the first thing to change.

CONTENTS

Five places AI pays first

Anum · AI enablement and change management · anum.com.au · ABN 80 697 129 525
Current at August 2026. Every legal, regulatory and pricing claim in this guide carries the source it came from. Anum is not a law firm and this is not legal advice.

Five places AI pays first

Five processes, what actually changes in each, and roughly what each one is costing while it stays the way it is. These five are examples. The same work covers agents, autonomous workflows, and whatever else a process turns out to need.

01 Quotes

Whoever prices your jobs records the details once, out loud, on the way back from site. The written quote comes back drafted in your own format, ready to check and send, instead of waiting for an evening.

> QUICK MATHS
> Writing a quote from scratch usually takes 15 to 20 minutes. Five quotes a week is over an hour of admin between jobs.

02 Reports

The report your team rebuilds every month out of the same four places gets assembled for them, with the figures pulled from the systems that already hold them and the commentary drafted from last month's.

> QUICK MATHS
> A monthly pack built by hand typically eats half a day. That is roughly six working days a year spent retyping numbers you already have.

03 Correspondence

Whoever answers your clients gets the answer with its source attached, found in your own documents and past emails rather than remembered or hunted for. The newest person on the team answers like the one who has been there longest.

> QUICK MATHS
> Two people interrupted three times a week for an answer that already exists in writing is over 300 interruptions a year.

04 Onboarding

A new starter asks your own material instead of asking the person beside them, and gets the page it came from. Weeks to competence rather than months.

> QUICK MATHS
> A new hire leaning on a colleague daily for two months has typically cost that colleague around 40 hours by the time they are independent.

05 Month end

The receipt, the invoice and the statement get read, coded and filed by the system. Whoever does your books approves rather than types.

> QUICK MATHS
> Finding out what you made 10 to 15 days after month end is roughly two working weeks, every month, spent waiting to know what already happened.

Every one of these is a set-up job before it is a build. None of them starts with buying software, and the honest answer is sometimes that a process should be left alone. The nine chapters that follow are the order we do it in.

1. What actually goes wrong

Nobody sets out to waste a year on this. The failures are boring and they repeat.

The licences are already paid for and nobody opens them

Microsoft 365 Copilot Business is AU$37.68 per user per month billed monthly, or AU$26.91 per user per month paid yearly, excluding GST. That yearly rate is a promotion running to 31 December 2026; Microsoft's standing price is AU$31.40. On 25 seats that is between AU$8,073 and AU$11,304 a year. If half those seats are never opened, the business is paying about AU$4,000 to AU$5,650 a year for nothing. That is not a hypothetical failure mode. It is the most common one we see, because the licences usually arrive as part of a bundle renewal rather than as a decision anyone made.

Microsoft Australia, Copilot pricing, checked 7 September 2026. Links and dates for every source in this guide are on the last page.

The work moves to personal accounts

When there is no approved tool and no written rule, people who want to get their work done use whatever is on their phone. A personal ChatGPT login. A free Claude account. A browser extension somebody found. The output is often good. The problem is that the business now has customer information sitting in accounts it does not own, cannot audit, and cannot get back if the person leaves.

The defaults on consumer plans are not the defaults on business plans, and the difference is the whole point. Anthropic states that on its commercial products it will not use your inputs or outputs to train its models by default. On its consumer plans, chats are used to improve models if you allow it, if a conversation is flagged for safety review, or if you opt in. Same company, same model, different plan, different answer.

Anthropic, "Is my data used for model training?", the commercial and the consumer answer.

The pilot never lands

Someone runs a trial. It works. Then it stops, because a trial has a finish line and a rollout does not. Nothing was written down, nobody was trained, the workflow around it never changed, and three months later the only person still using it is the person who ran the trial.

There is no owner and no number

Two failures that look separate and are the same failure. If nobody's job description contains this, it happens on top of somebody's real work until they get busy. And if nobody measured what the work cost before anything changed, nobody can say afterwards whether it was worth doing. The subscription renews on the assumption that it probably was.

Everything in the rest of this guide is aimed at those four.

2. Five questions to answer before you buy anything

Answer these on paper before a single licence is bought or a single tool is trialled. If you cannot answer them, you are not ready to buy, and buying will not make you ready.

1. Where does the week actually go?

Not where you think it goes. Sit with the three people who do the highest volume repetitive work in the business and watch them do it for an hour each. Write down every step, every copy and paste, every time they open a second system to look something up. The list you get from watching is different from the list you get from asking, and the difference is where the money is.

2. What is the cost of getting it wrong?

For each candidate task, ask what happens if the output is wrong and nobody catches it. A first draft of an internal email is one answer. A quote sent to a customer, a clinical note, a compliance return or a piece of advice is a completely different answer. This question decides how much human checking the workflow needs, and human checking is a real cost that has to go in the sums.

3. Does the Privacy Act apply to us, and what is in scope?

Most small businesses with an annual turnover of $3 million or less are not covered by the Privacy Act 1988. Plenty of small businesses are over that line, and some are covered regardless of turnover. The OAIC lists health service providers, businesses that trade in personal information, Commonwealth contract service providers, credit reporting bodies, residential tenancy database operators, reporting entities under the Anti-Money Laundering and Counter-Terrorism Financing Act 2006, Consumer Data Right accredited businesses and others. Annual turnover for this purpose means all income from all sources; it does not include assets held, capital gains or the proceeds of capital sales.

Work out which side of that line you are on before you decide what may go into a tool. If you are not covered, your contracts with your own customers may still bind you, and a breach still costs you the customer.

OAIC, "Small business".

4. Who owns this, and how much of their week do they get?

One named person, with authority, and a number of hours. The National AI Centre's guidance puts this first as well: assign a senior leader as the overall AI governance owner, someone with enough authority and enough understanding of the risks to oversee AI use across the organisation. In a small business this is almost never a new hire. It is usually an ops manager or a practice manager given four hours a week and the mandate to spend them.

National AI Centre, "Guidance for AI adoption: foundations", practice 1, Decide who is accountable.

5. What does the current version of this work cost, in a number?

Pick the measure before you change anything. Minutes per quote. Days from job finished to invoice sent. Number of documents that came back for a correction. Take the measurement for two weeks with nothing switched on. If you skip this step you will never be able to prove the result, and you will spend next year arguing about whether it worked.

3. What can and cannot go into a public model

This is the chapter people skip and it is the one that ends careers. Read it before you write the policy in chapter 4.

The regulator's position, plainly

The Office of the Australian Information Commissioner published guidance on privacy and the use of commercially available AI products in October 2024, updated in January 2025. Its recommendation is direct:

> As a matter of best practice, the OAIC recommends that organisations do not enter personal information, and particularly sensitive information, into publicly available AI chatbots and other publicly available generative AI tools, due to the significant and complex privacy risks involved.

OAIC, "Guidance on privacy and the use of commercially available AI products".

Three things in that sentence are load bearing. "Publicly available" is doing work: the guidance treats a tool inside your own tenancy, where information does not leave your control, differently from a public chatbot. "Personal information" is broader than most people assume. And "particularly sensitive information" points at a defined category that carries extra obligations.

Input is use, and sometimes it is disclosure

The OAIC's framing is the useful one. Putting personal information into an AI system is a **use** if the data stays within your organisation's control, and a **disclosure** if it is released from your effective control to someone outside the business. Either way APP 6 applies, which means the information can only be used or disclosed for the purpose it was collected for, unless an exception applies.

The guidance is blunt about how hard that exception is to rely on for AI:

> Given the significant privacy risks that may be posed by AI systems, and strong levels of community concern around the use of AI, in many cases it will be difficult to establish that a secondary use for AI-related purposes was within reasonable expectations.

Two more that catch people out. If an AI system generates or infers personal information, that is a **collection** and APP 3 applies to it. And consent for sensitive information cannot be implied merely because someone was notified.

A working list

This is the shape of the list we write with clients. Yours will differ, and the point of chapter 4 is that you write your own.

CATEGORY	PUBLIC MODEL (PERSONAL OR FREE ACCOUNT)	APPROVED BUSINESS TENANCY
Public information, marketing copy, published material	Yes	Yes
Internal drafts with no names and no client detail	Care	Yes
Customer names, addresses, contact details, account history	No	Only if the purpose matches why it was collected
Health information, and anything else defined as sensitive information	No	Only with consent, and usually not at all
Employee records, performance material, disciplinary matters	No	Restricted, named approvers only
Financial records, bank details, payment data	No	No
Anything under a confidentiality clause or an NDA	No	Read the clause first
Credentials, API keys, passwords	No	No
Third party material you licensed but do not own	No	Read the licence first

"Care" in that table means: strip the names, and if you cannot tell whether a document identifies someone, treat it as if it does.

Employee data has its own rule, and it is narrower than people think

Private sector employers get an exemption for acts and practices directly related to a current or former employment relationship and an employee record. It does not cover prospective employees, so a resume screening tool is not inside it. It does not cover volunteers. It does not cover tax file number information. If you are thinking of putting AI anywhere near hiring, do not assume the employee records exemption covers you.

OAIC, "Employee records exemption".

Where the data physically goes

APP 8 covers cross-border disclosure, and it is a real consideration when you choose a product: where the servers are, and whether personal information could be disclosed outside Australia. Ask the vendor. Get the answer in writing. If the answer is vague, that is the answer.

Read the vendor's own page for the exact plan you are on

Do not take a claim about a vendor from a blog, from a reseller, or from this guide. Read the vendor's page, for the specific plan you pay for.

Microsoft states, for Microsoft Copilot in its commercial services, that "prompts, responses, and data accessed through Microsoft Graph aren't used to train foundation LLMs, including those used by Microsoft Copilot."

Microsoft Learn, "Data, Privacy, and Security for Microsoft Copilot".

Anthropic states that by default it will not use your inputs or outputs from its commercial products to train its models, and separately describes a different default on its consumer plans. Both links are on the last page.

OpenAI publishes its business pricing openly and we quote it in chapter 1, but its enterprise privacy page would not load for us while writing this guide, so we are not going to characterise its position for you. Read it yourself at https://openai.com/enterprise-privacy before you put anything into a ChatGPT plan, and read it for the plan you actually hold.

The accounts that already exist

Before you write a policy, find out what is already in use. Run an amnesty: ask everyone, in writing, to list every AI tool they have used for work in the last three months, including free ones and browser extensions, with no consequences for answering honestly. You will get a longer list than you expect, and it is the only accurate starting point you will ever get. A policy written against the tools you assumed were in use protects nothing.

Then keep the list. The National AI Centre's fourth practice is to create and maintain an AI register: what systems you use, how you use them, and what matters about them, including AI embedded inside other software you already run. It publishes a free register template.

National AI Centre, AI systems register template.

How long it is kept, and who can see it

Two questions to put to every vendor before you approve a tool, in writing. How long is our content retained after we delete it? And which of your staff can read it, under what circumstances?

The answers are specific to the product and to the plan, and they are not always short. Anthropic, for example, states that where a user submits a thumbs up or thumbs down report, the entire related conversation is stored in its back end for up to five years. That is a reasonable arrangement, and it is also the sort of detail that never comes up in a demonstration. Ask, read the answer, and write it into the register above.

Anthropic, "Is my data used for model training?".

If something does go into the wrong place

If the Privacy Act covers you, the Notifiable Data Breaches scheme applies. An eligible data breach has three elements: unauthorised access to, or unauthorised disclosure of, personal information, or a loss of personal information; that is likely to result in serious harm to one or more individuals; and the entity has not been able to prevent the likely risk of serious harm with remedial action. Where you suspect one, you must take all reasonable steps to complete an assessment within 30 calendar days of becoming aware of the grounds for suspicion. If it is an eligible data breach, you notify affected individuals and the Commissioner as soon as practicable.

Thirty days is a maximum, not a target. Decide now who runs that assessment, because deciding during one is too late.

OAIC, "Part 4: Notifiable Data Breach (NDB) Scheme".

4. An acceptable use policy you can adapt

One page. If it runs to five, nobody reads it and it protects nobody.

Copy the skeleton below, fill in the bracketed parts, delete what does not apply, and put it where people will actually see it. The National AI Centre publishes a free AI policy template and an AI register template if you want a second structure to compare against.

National AI Centre, AI policy template and AI systems register template.

[BUSINESS NAME] acceptable use of AI tools

`Version 1.0. Owner: [NAME, ROLE]. Reviewed: [DATE]. Next review: [DATE + 6 months].`

1. Why this exists. We use AI tools to do our work faster and more consistently. This page says which tools are approved, what may and may not go into them, and who checks the output. It applies to everyone, including contractors.

2. Approved tools. Only these, on your work account: [TOOL, TOOL, TOOL]. Personal accounts are not approved for any work task, including free ones. If you want a tool that is not on this list, ask [NAME] before you use it, not after.

3. What must never go into any AI tool. Customer or client personal information. Health information and anything else that is sensitive information. Employee records, performance and disciplinary material. Bank details, payment data and financial records. Passwords, API keys and credentials. Anything covered by a confidentiality clause, an NDA or a client contract. Anything you would not be comfortable seeing outside this business.

4. What may go in, with care. Internal drafts once names and identifying detail have been removed. Published and public material. Your own notes. If you are not sure whether a document identifies someone, treat it as if it does.

5. Who checks the output. Nothing generated by an AI tool leaves this business without a named person reading it first. For [HIGHEST RISK OUTPUT: quotes, advice, clinical notes, compliance returns], that person is [NAME or ROLE]. AI output is a draft. You are responsible for what you send.

6. Say when it is AI. If a customer is dealing with an AI tool rather than a person, tell them. If AI materially shaped a document that goes to a client, and they would want to know, tell them.

7. Accuracy. These tools produce confident text that is sometimes wrong, including invented facts, invented citations and invented numbers. Check every figure, name, date and reference against the source. Do not use an AI tool as the source.

8. If something goes wrong. If you think information went somewhere it should not have, tell [NAME] the same day. You will not be in trouble for reporting it. You will be in trouble for hiding it. [NAME] will assess it and, where the Privacy Act applies to us, complete that assessment within 30 calendar days.

9. Changes. This page is reviewed every six months and whenever we add or drop a tool. [NAME] owns it.

Signed: _____ Date: _____

Three notes on adapting it. Keep clause 3 as a list of things, not a list of principles, because a principle can be argued with at 4:55pm and a list cannot. Name real people in clauses 5, 8 and 9, not departments. And set the review date when you sign it, not when you get around to it.

5. Choosing the first use case

The first thing you change decides whether there is a second thing. Pick it on evidence.

Score every candidate task out of 5 on four factors and multiply the scores. Multiplying rather than adding is deliberate: a task that scores 1 on any factor should fall to the bottom, and adding lets a strong score on one factor hide a fatal weakness on another.

FACTOR	1	3	5
Frequency. How often does this happen?	Monthly or less	Weekly	Many times a day
Time. How long does one instance take?	Under 5 minutes	15 to 30 minutes	Over an hour
Error cost. What does a mistake cost?	Real money or real trust	Some rework	Almost nothing
Checkability. How fast can a person tell it is right?	Needs an expert and an hour	A quick read by the right person	Obvious at a glance

Frequency and Time score the size of the prize. Error cost and Checkability are scored so that **low risk and easy to verify score high**, which is the opposite of how people instinctively rank the exciting use cases. That is the point. The exciting ones are usually the ones where a wrong answer is expensive and nobody can tell it is wrong.

A blank sheet to fill in:

TASK	FREQUENCY	TIME	ERROR COST	CHECKABILITY	SCORE

An illustration, with numbers you should replace with your own. Say a business scores "drafting the standard scope section of a quote" at frequency 4, time 3, error cost 3, checkability 5, for 180. It scores "deciding which supplier quote to accept" at frequency 2, time 4, error cost 1, checkability 1, for 8. The first one is the first use case. The second one is a conversation for next year, if ever.

Two rules that save people from themselves.

Start where the output is checked anyway. If a human already reads every one of these before it goes out, you get the speed without adding a new control.

Do not start with the thing the loudest person wants. Start with the thing that scores highest. If those are the same thing, good.

6. Rolling it out to people, not to software

Switching a tool on takes an afternoon. Getting people still using it in a month is the actual work, and it is where most of this fails.

Train on their job, not on the tool

A session about what the tool can do produces nodding and no change. A session where each person does their own real task, in the new way, with their own real files open, produces change. Bring the work into the room. If the training does not include the person completing a piece of their own work, it is a demonstration, not training.

Write the workflow down, then delete the old one

If both paths still exist, the old path wins the first busy week. Retire the old template. Move the old form. Change the checklist. The new way has to be the easy way, and easy usually means "the only one still on the intranet".

Name the checker before the first output

Chapter 4, clause 5. Someone reads it before it leaves. Decide who, in writing, before anyone generates anything.

Build the prompt library as you go

Every time someone works out a form of words that reliably produces the right output, it goes in a shared document with the task name on it. This is the cheapest asset in the whole programme and almost nobody builds it. Six months in, it is the difference between one person who is good at this and a team that is.

Two weeks in, sit with them again

Not a survey. Sit next to three people and watch them do the task. You will find out in twenty minutes what a survey will not tell you in a month: which step they skip, which bit they gave up on, and the workaround somebody invented that should become the standard.

What "done" means

Write it down before you start. We use this shape, and you can steal it: **done is when the named workflow is the only documented way to do the task, the people who do it have each completed it unassisted, the checker is named, the policy covers the tools in use, and the measure from chapter 7 has been taken twice.**

Anything short of that is in progress, whatever the slide says.

7. Measuring whether it paid

The rule is one sentence: same measure, taken before and taken after, by the same method.

Everything else is detail, and the detail is where people cheat without meaning to.

Take the baseline before you change anything

This is the step that gets skipped, and skipping it is unrecoverable. You cannot reconstruct a baseline later. Two weeks of the real measure, taken while the old way is still the only way.

Pick a measure that survives an argument

Good measures are counts and durations that somebody else could take and get the same answer.

- Minutes from job complete to invoice sent
- Number of quotes issued per person per week
- Days to close the month
- Documents returned for correction, as a share of documents sent
- Time from enquiry received to first substantive reply

Bad measures are ones you cannot recount next quarter: "the team feels faster", "we saved about a day a week", "everyone loves it". Sentiment is worth collecting and is not evidence of a result.

Count the whole cost

The saving is not the gross time saved. It is time saved, minus the time spent checking the output, minus the licence cost, minus the time the owner spends running it. Write all four down. A use case that saves 6 hours a week and costs 4 hours a week of checking is a 2 hour result, and it might still be worth doing, but you should know that it is a 2 hour result.

Take it twice

Once at the end of the rollout, once thirty days after the rollout finishes. The gap between those two numbers is the only honest measure of whether the change stuck, and it is the number most programmes never collect. A result that is real at week 8 and gone at week 12 was not a result.

Say what it did not do

If two of your five use cases paid and three did not, the report says that. A measurement regime that only ever produces good news is not a measurement regime.

8. When to stop, and what "not yet" looks like

Sometimes the answer is that this is not worth doing, or not worth doing here, or not worth doing this year. That is a real finding and it should be available to you at every step. Here is what it looks like.

The task is not repeatable. If every instance is different in a way that requires judgement, there is nothing to standardise and the tool has nothing to learn from. Fix the process first, or leave it alone.

The underlying data is a mess. AI over documents only works if the documents are findable, current and correct. If your file share has four versions of the fee schedule and nobody knows which is live, the honest first project is cleaning that up, and it is not an AI project.

Nobody has the hours. If nobody in the business can commit real weekly time to owning it, do not start. An unowned rollout is a licence bill with training attached.

A wrong answer is expensive and nobody can check it fast. Both together, on the same task, is a stop signal. Either one alone is manageable.

A rule already forbids it. A client contract, a professional obligation, a funding agreement or a regulator's position can put a task out of scope regardless of how good the tool is. Read the clause. This is what chapter 2, question 3 was for.

The prize is too small to survive the effort. If the whole task is 20 minutes a month, you will spend more managing the change than you will ever get back.

There is a version of this that is a delay rather than a refusal, and it is common. The work is right, the tool is right, and the business is in the middle of an ERP migration or a merger or a busy season. Adding a change programme on top of that is how both fail. Write it down, put a date on it, and come back.

Saying no to a use case is not the same as saying no to the whole thing. It is usually the fastest route to the one that works.

9. Where this goes next

You can run every chapter of this guide yourself. Plenty of businesses do, and if you have someone with the hours and the temperament for it, that is a better outcome than hiring anyone.

If you would rather not, that is what Anum does. We give a business the AI function it cannot justify hiring: discovery, prioritisation, governance and policy, implementation, staff training, measurement, and an ongoing fractional AI lead if it is useful afterwards.

Discovery & Training

The first thing anyone pays for. The length and the price are agreed in writing before it starts, because both depend on the size of the business and how many teams are in scope.

You end up with the same artefacts this guide describes, built from your actual work rather than from a template:

- A map of where the week goes, taken from watching the work, not from a survey
- A scored list of candidate use cases, using the method in chapter 5
- The baseline measure from chapter 7, taken before anything changes
- A written AI use policy and data handling rules, adapted to your obligations
- A training session with your team, built from what we found: getting more out of the AI you already pay for, and using it safely
- A proposal for what to do next, scoped and priced

If Discovery leads to a Rollout and you proceed within 60 days, we credit the Discovery fee towards the implementation. If you do not proceed, everything produced through Discovery is still yours, and there is nothing else to pay.

The part we will say out loud

If Discovery does not find at least one use case with a clear payback inside six months, we will tell you the Rollout is not worth doing yet, and we refund the fee. Chapter 8 is on this page because we mean it.

We will not put AI into a workflow that should not have it.

Talk to us

A 30 minute call, free, nothing to prepare. Tell us where your team is losing time. You do not need to know what the solution looks like and you do not need to have decided anything.

Every message through the site is read by the person who would do the work, and you get a reply within two business days either way.

Let's work through what's costing you the most and turn it into revenue

`anum.com.au/#contact · hello@anum.com.au`

Where the facts in this guide came from

Every claim about law, regulation or price above is attributed where it is made. The links are collected here, so the argument reads without them:

- OAIC, Guidance on privacy and the use of commercially available AI products: https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/guidance-on-privacy-and-the-use-of-commercially-available-ai-products
- OAIC, Australian Privacy Principles quick reference: https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-quick-reference
- OAIC, Small business: https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/organisations/small-business
- OAIC, Employee records exemption: https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/organisations/employee-records-exemption
- OAIC, Part 4: Notifiable Data Breach (NDB) scheme: https://www.oaic.gov.au/privacy/notifiable-data-breaches/preventing-preparing-for-and-responding-to-data-breaches/data-breach-preparation-and-response/part-4-notifiable-data-breach-ndb-scheme
- National AI Centre, Guidance for AI adoption: foundations: https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations
- Microsoft Learn, Data, Privacy, and Security for Microsoft Copilot: https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-privacy
- Anthropic, Is my data used for model training? (commercial): https://privacy.claude.com/en/articles/7996868-is-my-data-used-for-model-training
- Anthropic, Is my data used for model training? (consumer): https://privacy.claude.com/en/articles/10023580-is-my-data-used-for-model-training
- National AI Centre, AI policy template: https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/create-ai-policy
- National AI Centre, AI systems register template: https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/ai-systems-register
- Microsoft 365 Copilot pricing (Australia): https://www.microsoft.com/en-au/microsoft-365-copilot/business
- Anthropic, Claude pricing: https://claude.com/pricing
- OpenAI, ChatGPT business pricing: https://openai.com/business/chatgpt-pricing/
- OpenAI, enterprise privacy: https://openai.com/enterprise-privacy

All sources above were checked on 24 August 2026, and the Microsoft pricing page was rechecked on 7 September 2026. Prices are the vendors' own published figures, in the currency the vendor publishes them in, and are not converted. Claude Team is US$20 per seat per month billed annually and US$25 billed monthly. ChatGPT Business is published at A$28 per user per month billed annually, minimum two users; only ChatGPT Enterprise is priced on application.

Anum is a Melbourne technology company. ABN 80 697 129 525. This guide is general information about how to run an AI adoption programme. It is not legal advice and it is not a privacy compliance assessment. Where a decision turns on your obligations, read the source, and get advice from someone qualified to give it.