



WORKING COPY

THE PROFESSIONAL AI USE PROTOCOL

*The one-page Protocol from **You're Already Using AI at Work**, with room to write*

The book builds one page, one line from each chapter. This working copy is the same Protocol, set out so you can fill it in as you read. Each line has the questions to answer, the tables from its chapter, and space to write. The last page is the one page itself: copy your final answers onto it and keep it where you'll see it.

How to use it

- Work through a line after you finish its chapter. Each takes five to fifteen minutes, because the thinking was done in the chapter.
- Write in pencil. Every line except the first will change, and the review on Line 7 is how you keep up.
- Date everything. The Protocol is a set of dated claims about the world, and dated claims go off.
- Never copy a prompt itself into this document. Record what it was for, not what it said.
- Keep the finished page to one page. If it spills onto a second, it's become a policy document.

What's inside

Line	From	What it records
1. The three questions	Chapter 1	Whose is it? What if it's wrong? Who needs to know? Plus the ten-prompt audit.
2. My data classes	Chapter 2	My strictest class, and my moves for getting tasks down a class.
3. My setup	Chapter 3	Each tool: account, plan, date terms were checked, classes allowed.
4. My check	Chapter 4	Source, scope, sense, sign, plus my own highest-risk error type.
5. My disclosure position	Chapter 5	Default, triggers, words; date of the afternoon; who to ring.
6. My shortlist	Chapter 6	Three to five tasks, each with class, rung and check. Plus my stop list.
7. My review	Chapter 7	When I next review, and what triggers an early one. Plus the quarterly review checklist.
The one page		Your finished Protocol.

My name _____ Started on _____

This working copy gives general information, not legal, medical or other professional advice. Rules differ by profession and jurisdiction and change over time, so check the current position with your regulator, employer and indemnity insurer before relying on anything here.
© 2026 J.D. Ashby. Print as many copies as you need for your own use and your team's. Version 1, October 2026.

THE THREE QUESTIONS

From Chapter 1: Audit What You've Already Sent

Every prompt you have ever sent at work can be judged by three questions. They track the three commonest ways a professional gets into trouble with AI: sending something that wasn't theirs to send, relying on something that was wrong, and keeping quiet when someone needed to know.

LINE 1 IS ALREADY WRITTEN. ASK IT BEFORE YOU PRESS ENTER.

WHOSE IS IT? WHAT IF IT'S WRONG? WHO NEEDS TO KNOW?

1. Whose information was that?

If it's yours alone, or already published, relax. If it's someone else's, you were holding it on their behalf. Removing the name is not the same as removing the person: a date of birth, a rare diagnosis in a small town, a matter number or a site address can do the same job. **Shoulder test:** if someone from the same professional or local circles as the person read your prompt over your shoulder, could they tell who it's about?

2. What happens if it was wrong?

Assume the output is sometimes wrong, and ask where that wrongness would have ended up. Binned, it went nowhere. Checked against a source you trust, you had a fair chance to catch it. Sent out under your name with no independent check, any error in it now belongs to you.

3. Who needed to know you used it?

Sometimes nobody. But the patient or client, your employer, a regulator, a court, a journal or a publisher may have had a right, or a reasonable expectation, to know a machine was involved. **Test:** would anyone be surprised, or annoyed, to find out this piece of work went through an AI tool?

Scoring a prompt

Question	Green	Amber	Red
1. Whose information?	Yours alone, or already public	Someone else's, but genuinely stripped of anything that identifies them	Someone else's, and they could be identified from it
2. What if it was wrong?	Nothing left your desk	You checked it against a trusted source before using it	You used it without an independent check
3. Who needed to know?	Nobody plausibly would	You're not sure	Someone had a right or expectation to know, and wasn't told

A prompt takes the colour of its worst answer. Two greens and a red is a red. Whether a red is a breach in law is a question for the people on your Line 5.

THE TEN-PROMPT AUDIT

From Chapter 1. About twenty minutes. Repeat it at every quarterly review (Line 7).

1. **Find everything.** Open the history in every AI tool you've used for work. Include the ones people forget: the app on your phone, the personal account, the AI button in your email or word processor, the transcription or scribe tool, the browser extension. If a tool keeps no history, reconstruct the last fortnight as best you can.
2. **Pick the last ten work prompts.** Not the ten best, not the ten worst. The last ten, in order.
3. **Write one line for each, saying what it was for.** Record the purpose, not the content. Don't copy the prompt here, or you've made another copy of the information you're worried about.
4. **Score each one** against the three questions (G, A or R), then give it its overall colour: the worst of the three.
5. **Count** the greens, ambers and reds.

Audit date _____ Covers prompts from _____ to _____

#	What it was for	Tool and account	Q1 Whose?	Q2 If wrong?	Q3 Who knows?	Overall
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						

Greens _____ Ambers _____ Reds _____

The "tool and account" column matters more than it looks. The same prompt sent from your employer's enterprise account and from a free app on your phone are two different things. Chapter 3 explains why.

What to do with a red

- ☐ **Find out where it went:** which product, which account, what the settings were. Check the product's settings and the provider's published terms. Don't guess.
- ☐ **Know what deleting does.** It removes the conversation from your view. Whether it removes the provider's copy is in their terms.
- ☐ **Write it down:** date, tool and account, the kind of information (not the information itself), and what you've done since.
- ☐ **Don't decide alone** whether it needs reporting. Ring the people on Line 5.
- ☐ **Fix the habit, not just the prompt.** Find the pattern and pause it until you've done Lines 2 and 3.

Where each amber goes

Not sure the information was really de-identified? **Line 2.**

Not sure your tool or account was right for it? **Line 3.**

Checked it, but not against the right source? **Line 4.**

Not sure who needed telling? **Line 5.**

Wondering whether the task was worth handing to a machine at all? **Line 6.**

MY DATA CLASSES

From Chapter 2: Sort What You're Allowed to Put In

Class	What it covers	Default
A: Public	Already published, with nobody in particular in it (guidelines, legislation, standards, articles), or a general question.	Green
B: Yours	Your own thinking and material, with nobody else's information in it: a talk, a template with blanks, your CV.	Green
C: Your organisation's	Not about any individual, but not yours: internal figures, unpublished policies, rosters, budgets, tenders, a client's business information. For lawyers and accountants, client-matter facts stay Class C even de-identified.	Amber
D: Someone's, de-identified	About a person, but <i>genuinely</i> stripped of anything that could identify them.	Amber
E: Someone's, identifiable	About a person who could be identified from it, by name or by the combination of details. Health, mental health, sexual and reproductive health, criminal history, children, privileged and financial information weigh more.	Red

The strictest class in a prompt sets its risk. The letters sort; they don't rank. "Strictest" means the strictest row in Line 3's table. A prompt that is ninety per cent public guideline and one line of identifiable patient detail is a Class E prompt.

Direct identifiers: take these out

- ☐ Names, including relatives, carers, witnesses and business partners
- ☐ Dates of birth; exact dates of admission, appointment, settlement, inspection or incident
- ☐ Addresses, postcodes and suburbs, including a site or property address
- ☐ Phone numbers and email addresses
- ☐ Record, matter, project, job, tax, policy and account numbers
- ☐ Photographs, scans, drawings with a title block, and signatures
- ☐ Letterhead and file headers

The four moves

Remove what the task doesn't need.

Generalise: 47 becomes "late forties", a named town becomes "a small regional town".

Substitute placeholders: Patient A, Client B, [NAME], [DATE]. Fill the real details in afterwards, in your own system. For privileged material, abstract instead.

Abstract: ask about the problem, not the case. An abstract question is often Class A.

Indirect identifiers: watch the combination

- ☐ Rare conditions, unusual injuries or circumstances
- ☐ Occupation, especially a distinctive one
- ☐ Small places, or a place plus almost anything else
- ☐ Exact ages, particularly at the extremes
- ☐ Family structure
- ☐ Events that were public, or gossiped about
- ☐ Distinctive assets, buildings or businesses
- ☐ The specific facts of a legal matter, as identifying as a name to anyone who knows the case

The two tests

Shoulder test: if someone from the same professional or local circles as the person read your prompt over your shoulder, could they tell who it's about?

Colleague test: if someone in your own workplace read it, could they match it to a file?

If either answer is yes, it's still Class E, whatever you've removed. Watch whole documents, screenshots and audio: they are hard or impossible to de-identify.

MY TASK MAP

Start with the ten tasks from your audit. Add the jobs you'd like to use AI for, or quietly stopped. Aim for fifteen to twenty.

#	Task, in a few words	Class as done	Lowest possible class	Move that gets it there
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				
11				
12				
13				
14				
15				

Leave the lowest class blank if you're not sure. The gap between the two class columns is the whole opportunity.

QUESTIONS TO ANSWER

Of the tasks you intend to keep doing, which has the strictest class in the right-hand column? Which two or three tasks gain most from moving down a class, and which move does it?

LINE 2: MY DATA CLASSES

My strictest class _____

My moves, e.g. "Referral letters: placeholders for name and identifiers, fill in locally."

MY SETUP

From Chapter 3: Choose a Setup That Matches Your Risk, Not Your Budget. Answer the six questions once per tool, and date them.

The rungs

Rung 1: Consumer accounts. The free app, or a personal plan, paid or unpaid. Paying for a personal plan is not the same as having a business agreement.

Rung 2: Business, enterprise and API. Sold to organisations, with a contract attached. It usually means your employer chose and negotiated it. It's a contract, not a colour scheme.

Rung 3: Local models. Runs entirely on your own computer, so the prompt never leaves the machine. You pay in capability, effort and your own security, and your employer's rules still apply.

The rung you didn't know you were on: the AI button in your email, the word processor assistant, the browser extension, the scribe. Ask the six questions of these too.

Minimum setup for each class

Class in the prompt	Minimum setup
A: Public	Any rung
B: Yours	Any rung
C: Organisation's (and, for lawyers and accountants, any client-matter facts, even de-identified)	Rung 2 approved by the organisation or client that owns the information, or Rung 3
D: De-identified	Rung 2 by default. Rung 1 only for prompts that are close to abstract, with model training switched off and a temporary or incognito chat where the product offers one, and only where your employer hasn't said otherwise.
E: Identifiable	Rung 2, only where your organisation has approved that specific tool for that kind of information under an agreement. Or Rung 3, within your employer's rules.

Where a prompt holds more than one class, the strictest row applies. This is the minimum, not the ideal, and your employer's rules sit on top of it. Treat a scribe as Class E, because audio is so hard to de-identify.

Reading the terms

- ☐ **Find the page for your plan,** not the first search result.
- ☐ **Look for the answers, not the adjectives.** "Conversations are retained for 30 days" is an answer.
- ☐ **Watch the verbs.** "We will not" is a commitment. "We may" is a permission slip the provider has written for itself.
- ☐ **Save it and date it:** the URL, the date you read it and the key answers.
- ☐ **Ask your employer first** which classes of information it's approved for. If nobody can answer, that's an answer too.

MY TOOLS

One card per tool. Start with the one you use most.

Tool

Account and plan

Rung

Terms page (URL)

Date checked

Question	Answer, in the provider's own words where you can
1. Is it kept?	
2. Is it trained on? On by default? Can I turn it off?	
3. For how long? What happens when I delete it?	
4. Who can see it?	
5. Where is it processed?	
6. Is there a contract that says so, with me or my employer?	

Approved by my employer for

Classes I'll put into it

Tool

Account and plan

Rung

Terms page (URL)

Date checked

Question	Answer, in the provider's own words where you can
1. Is it kept?	
2. Is it trained on? On by default? Can I turn it off?	
3. For how long? What happens when I delete it?	
4. Who can see it?	
5. Where is it processed?	
6. Is there a contract that says so, with me or my employer?	

Approved by my employer for

Classes I'll put into it

More tools? Print this page again.

LINE 3: MY SETUP

One line per tool: tool, account and plan; date terms checked; classes allowed. E.g. "Personal ChatGPT, training off, checked 5 October 2026. Class A and B only."

MY CHECK

From Chapter 4: Get Output You'd Put Your Name On

How hard to check

Where the output goes	Check
Nothing leaves your desk (brainstorming, the shape of an argument)	A quick sanity read
It informs a decision you then make yourself	Check the parts the decision rests on
It goes out under your name: a record, letter, advice, filing or publication	The full routine. Every time. Including Friday afternoons.

The routine: four checks, in this order

Check	Ask	Done
1. Source	Have I traced every reference, dose, threshold, statute, standard, date and number to a source I opened myself? A citation the model gave me is a lead, not a source.	
2. Scope	Is it for the right jurisdiction, the right version or date, the right population and the right facts? Does it answer the question I actually asked?	
3. Sense	Read as a sceptical colleague: does it add up, do the numbers sum, does anything contradict itself? What would I expect in a good answer that isn't here?	
4. Sign	Would I sign this if the AI didn't exist? Any sentence I can't defend: check it or cut it.	

My highest-risk error type

Tick the one most likely to bite in your work.

- ☐ **Fabricated:** references, cases, quotes or statistics that don't exist
- ☐ **Outdated:** a changed threshold, an amended section, a withdrawn drug
- ☐ **Wrong jurisdiction:** correct somewhere, just not where you are
- ☐ **Omission:** the one thing that matters is missing
- ☐ **Confident arithmetic:** sums, conversions, calculations
- ☐ **Scope drift:** a perfect answer to a slightly different question

Making checking cheaper

- ☐ **Give it the sources.** Paste the guideline and ask it to answer only from that document, quoting the passage for each point.
- ☐ **Ask for its assumptions.** End the prompt with: "List any assumptions you've made, and anything you're unsure about."
- ☐ **Use it for structure, not facts.** Supply the facts; let it supply the prose.
- ☐ **Ask for citations, then treat them as leads.**

LINE 4: MY CHECK

Source, scope, sense, sign, for anything that leaves my desk under my name. Plus my highest-risk error type, in words short enough to remember on a Friday afternoon.

MY DISCLOSURE POSITION

From Chapter 5: Settle What You Have to Declare. Block out two to three hours. Head a page with today's date.

Disclosure obligations come from four possible sources: the patient or client, the employer, the regulator, and anyone you submit work to. The afternoon checks all four, plus your insurer.

The afternoon	What I found, with the date ("no policy found" counts)
1. Read your own paperwork. Consent forms, engagement letters, privacy policy, information sheets. What do they say about AI, recording or third-party services?	
2. Ask your employer. Search the intranet for "AI", "artificial intelligence", and the names of tools I use. Email one question: which tools are approved, for what, and do we need to tell anyone?	
3. Search your regulator and professional body. Search each website for "artificial intelligence". Save anything relevant, with the date.	
4. Check your indemnity insurer. Look for AI guidance for members. Note the advisory line.	
5. List your submission targets. The courts, journals, funders or publishers I submit to. Check each one's current rules.	
6. Name who you'll ring about a red.	See "Who to ring" below.

QUESTIONS TO ANSWER

What do I tell people routinely, without being asked? In which situations do I go further? What exactly will I say when someone asks? (It should be true, short and calm. People rarely object to a tool. They object to finding out by accident.)

If you use a scribe: before the scribe starts, tell the patient what it does and ask if they're happy for it to run; note their answer; switch it off if they'd rather not. Recording a conversation without consent can be a criminal offence in some states and territories, so ask before you start, every time.

LINE 5: MY DISCLOSURE POSITION

My default _____

My triggers _____

My words _____

Date of my afternoon _____

Who to ring

Privacy officer _____ **Number** _____

Insurer's advice line _____ **Number** _____

Professional body's advisory service _____ **Number** _____

MY SHORTLIST

From Chapter 6: Pick the Work That's Worth It

AI pays off where checking is cheaper than doing, and costs you where it isn't. A task is worth handing over when the cost of doing it with AI (prompting, de-identifying, checking and fixing) plus the expected cost of an error that gets through is *clearly* less than the cost of doing it yourself. You'll underestimate the checking. Everyone does.

Usually wins: you already know the content, and the machine only changes its shape. Drafting from your own notes; changing tone or format; first drafts of standard documents; explaining public material you can verify; brainstorming; structure.

Usually loses: the checking amounts to redoing the work, or an error is expensive and hard to see. Facts you can't check quickly; summarising long documents you haven't read; novel judgement; final text in high-stakes documents; work where the thinking is the point.

Run your task map from Line 2 through four filters. **Class:** what's the lowest class? **Rung:** is the right setup available to me (Line 3)? **Check:** how long will the check take, honestly (Line 4)? **Worth:** is it clearly cheaper, once errors are counted? Then decide: **Keep, Redesign** (a different class, rung or way of asking) or **Stop**.

Task	Lowest class	Right rung available?	Honest check time	Worth it?	Decision

LINE 6: MY SHORTLIST

Three to five tasks, each with its class, rung and check. Three to five, done well, beats twenty done anxiously.

MY STOP LIST

The tasks I've decided not to hand over, and why. You'll be glad of it when a shiny new tool turns up.

MY REVIEW

From Chapter 7: Make It a System, Not a Stunt

Every line except the first depends on something outside your control, and Line 3 expires fastest. A one-off effort decays. A protocol you review compounds.

LINE 7: MY REVIEW

Next quarterly review

In my calendar as a recurring 45-minute appointment

Review the relevant line straight away, without waiting for the quarter, when:

☐

I start using a new tool, including an embedded one that just appeared (Line 3)

☐

a provider announces new terms, or I notice a new setting (Line 3)

☐

my regulator, employer, insurer or a submission target publishes anything about AI (Line 5)

☐

my role changes (Lines 2 and 6)

☐

something goes wrong (all of them)

Quarterly review checklist

Step	Done	Notes
1. Re-run the audit. Score my last ten work prompts on a fresh copy of the audit page. Does what I do match what my Protocol says?		
2. Re-check the dates. Open the terms page for every tool on Line 3. Has anything changed since the date I recorded? Update it, with the new date.		
Check my employer’s and regulator’s pages from Line 5, and my insurer’s guidance. Anything new?		
Has any tool been renamed, merged, given a new AI feature, or switched on without my asking?		
3. Re-sort the shortlist. Given steps 1 and 2, does anything on Line 6 need adding, stopping or redesigning?		
Update the one page, date it, and book the next review.		

When something goes wrong

☐

Find out what happened: which tool, which account, which settings, what information.

☐

Know what deleting does, and doesn't do.

☐

Write it down, with dates, at the time.

☐

Don't decide alone whether it needs reporting. Ring the people on Line 5.

☐

Fix the system, not just the incident. Which line of the Protocol failed? Update it.

MY PROFESSIONAL AI USE PROTOCOL

Name _____ Last updated _____ Next review _____

1. THE THREE QUESTIONS Asked before every prompt

WHOSE IS IT? WHAT IF IT'S WRONG? WHO NEEDS TO KNOW?

2. MY DATA CLASSES My strictest class, and my moves for getting tasks down a class

3. MY SETUP Each tool: account, plan, date terms were checked, classes allowed

4. MY CHECK Source, scope, sense, sign, plus my own highest-risk error type

5. MY DISCLOSURE POSITION Default, triggers, words; date of the afternoon; who to ring

6. MY SHORTLIST Three to five tasks, each with class, rung and check. Plus my stop list

7. MY REVIEW When I next review, and what triggers an early one

From *You're Already Using AI at Work* by Dr J.D. Ashby. J.D. Ashby is the pen name of a doctor registered with Ahpra and practising in Australia at the time of writing (2026). General information, not legal, medical or other professional advice.