




---

ENCUBIQ FIELD GUIDE

· AI STRATEGY PRACTICE

# Getting Real Work Out of AI

---

*A practical guide to Claude, ChatGPT, Perplexity, Microsoft Copilot and Google Gemini, for owners and operators of small and mid-sized businesses.*

## [ Contents

|                                         |    |
|-----------------------------------------|----|
| How to use this guide                   | 3  |
| 1. The five-minute orientation          | 4  |
| 2. Choosing your stack                  | 5  |
| 3. Choosing the right model for the job | 11 |
| 4. The anatomy of a prompt that works   | 12 |
| 5. Eight prompts, rewritten             | 14 |
| 6. Five habits that compound            | 19 |
| 7. Where it goes wrong                  | 20 |
| 8. Twelve jobs worth trying this month  | 24 |
| 9. The rules of the road                | 25 |
| 10. Your first thirty days              | 28 |
| Appendix A: Templates to copy           | 30 |
| References and disclosures              | 31 |
| About EncubIQ Consulting                | 32 |

### BEFORE YOU START

Nothing in this guide requires a technical background, a developer, or a budget beyond a few subscriptions. Every technique here can be tested in an afternoon.

## [ How to use this guide

Most writing about artificial intelligence is aimed either at engineers or at nobody in particular. This guide is aimed at the people who carry the business: the owner who answers the phone and still has to get the newsletter out on Friday, the general manager or department head who has eleven people asking whether they are allowed to use any of this yet, and the CFO or COO at a company where approving the spend, writing the policy and answering for both land on the same desk.

Those readers need different things from the same material. If you do the work yourself, sections 4 through 7 will change your output this week. If you are responsible for other people doing the work, or for what it costs and what it exposes, read those, then treat sections 9 and 10 as the real content: the trust boundary, the eight-line policy, and a rollout order that does not leave you retrofitting governance onto an established habit.

The argument it makes is narrow and easy to test. Most of the disappointment people report with these tools traces back to two fixable mistakes. The first is using the wrong tool for the job. The second is typing a one-line request where a six-line brief was needed. Fix both and the tools stop feeling like a novelty and start behaving like a capable assistant who happens to work at three in the morning.

The numbers below are worth holding lightly. They come from credible surveys, but they are surveys.



Those figures come from the US Chamber of Commerce, an IDC study commissioned by Microsoft, and a Protiviti and London School of Economics study respectively. Set against them, the US Census Bureau, which asks a stricter question about production use rather than any use at all, put adoption among businesses under 250 employees at 8.8% in August 2025. Both pictures are true. Plenty of businesses have opened the app; far fewer have wired it into how work actually gets done. This guide is about the gap between those two states.

## [ What this guide is

- A working comparison of the five assistants most small businesses will actually encounter, with honest notes on what each one is bad at.
- A method for writing requests that produce usable output on the first or second attempt rather than the fifth.
- Worked before-and-after examples you can adapt to your own business this week.
- The governance basics: what you can safely paste in, what you cannot, and what you now have to disclose.

## [ What it is not

- A prediction about where any of this is going.
- A recommendation to replace anyone on your team.
- A technical manual. There is no code here, and none is needed.

# 1 The five-minute orientation

If you have used one of these tools and come away unimpressed, it is worth knowing what changed. The versions that reached the public in 2022 and 2023 were text predictors with a chat window bolted on. What you are buying in 2026 is closer to a workspace: it can read your files, search the live web, run calculations, hold a project's context across weeks, and produce a working document rather than a wall of text.

## [ 1.1 What you are actually buying

Strip away the branding and all five products are the same three things in different proportions. There is a language model that generates text. There is a set of tools it can reach for, such as web search, a code sandbox for crunching a spreadsheet, or a connection to your Google Drive or SharePoint. And there is a memory layer that decides how much of your prior context it carries into the next request.

The practical consequence is that the assistant knows nothing about your business until you tell it or connect it. It does not know your margins, your tone of voice, your best customer, or why last October was a disaster. Almost every complaint about generic output comes back to this. You are asking a capable stranger to write about a company they have never heard of.

### THE CORE IDEA

Treat the assistant as a bright new hire on their first morning. Talented, fast, entirely without context. Anthropic uses this exact framing in its own guidance, and it is the single most useful mental model in this guide.

## [ 1.2 Context windows, in plain terms

A context window is how much material the assistant can hold in view at once: your instructions, the documents you upload, and the conversation so far. This is the number that decides whether you can drop in a sixty-page supplier contract and ask sensible questions about clause 14. The larger models now hold up to a million tokens, which is roughly 750,000 words, or a stack of documents you would need a weekend to read.

Two things follow. First, uploading the real document nearly always beats describing it. Second, very long conversations degrade. When a thread starts drifting or contradicting itself, start a fresh one and

paste in a short summary of where you had got to.

### [ 1.3 What it is genuinely good at, and where it still fails

| Reliably good                                                         | Still unreliable                                                                    |
|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| First drafts of almost anything, from a job posting to a client email | Arithmetic and data accuracy unless it is given a calculator or a file to work from |
| Summarising long documents, transcripts and email threads             | Anything requiring guaranteed correctness with no human check                       |
| Reformatting and restructuring material you already have              | Current facts, unless web search is switched on                                     |
| Explaining an unfamiliar concept at whatever level you ask for        | Knowing what it does not know, which is why it invents confidently                  |
| Extracting structure from messy notes, forms and free text            | Genuine originality; it produces the average of what has been written before        |
| Adapting tone for a different audience                                | Judgement calls that depend on context it has not been given                        |

Read that table as a division of labour rather than a scorecard. The left column gets you to a competent eighty per cent in a fraction of the time. The last twenty per cent, the part where judgement and accountability live, stays with you.

## 2 Choosing your stack



You do not need five subscriptions. Most businesses under a hundred people are well served by one paid workhorse for daily writing and analysis, plus one research tool that cites its sources. Everything else is a preference.

## [ 2.1 The short answer

If your business already runs on Microsoft 365, start with Microsoft Copilot; if it runs on Google Workspace, start with Gemini. Each has the same advantage: it works inside the documents and mailboxes you already have, on your own company data. If you are on neither, pick ChatGPT or Claude on merit: ChatGPT if you want the widest range of capabilities in one place, Claude if the majority of your use is writing and reading long documents. Add Perplexity for research regardless of which you choose, because the ability to click through to the source changes how much you can trust an answer.

## [ 2.2 Claude, from Anthropic

Claude has a reputation among people who write for a living as the best drafter of the group, and it holds up. It is less prone to the breezy listicle voice, it follows detailed style instructions closely, and it handles very long documents without losing the thread halfway through.

| Claude            |                                                                                                                                                                                                                                                                                   |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Model tiers       | Haiku for fast, high-volume work. Sonnet as the everyday default. Opus and above for hard reasoning. Version numbers move; the tier names have been stable.                                                                                                                       |
| Price             | Free tier, which now runs the same mid-tier model as Pro. Pro US\$20 a month, or US\$200 a year. Max at US\$100 or US\$200 a month for heavy users, monthly billing only. Team from US\$25 per seat a month, or US\$20 billed annually, minimum five seats. Enterprise is quoted. |
| Standout features | Projects, which hold your files and standing instructions in one workspace. Artifacts, which render a working document, calculator or chart beside the chat so you can iterate on it. Very large context. Connectors to Drive, Gmail, Slack and Notion.                           |
| Best for          | Proposals, client correspondence, contract and transcript review, policy documents, anything where the writing itself is the deliverable.                                                                                                                                         |
| Watch out for     | Consumer plans now train on your conversations unless you opt out. See section 9.1 before you paste anything client-related.                                                                                                                                                      |

## [ 2.3 ChatGPT, from OpenAI

The generalist, and the safest single choice if you want one tool for a small team with mixed needs. It writes well, analyses spreadsheets without you writing a formula, generates images, holds a voice conversation, and now routes your question automatically to a fast or a slow model depending on how hard it is.

| ChatGPT           |                                                                                                                                                                                                                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Model tiers       | A fast default for everyday chat, a Thinking mode for harder problems, and a research-grade tier on paid plans. In most cases the router picks for you.                                                                                                                         |
| Price             | Free tier, ad-supported in the United States. Go US\$8 a month. Plus US\$20, monthly only, and the sensible default for an individual. Pro at US\$100 or US\$200 a month. Business US\$25 per seat a month, or US\$20 billed annually, minimum two seats. Enterprise is quoted. |
| Standout features | Advanced data analysis, where you upload a spreadsheet and ask questions in English. Deep Research, which reads dozens of sources and returns a cited report. Projects, Canvas for side-by-side editing, image generation, and a wide set of connectors.                        |
| Best for          | Mixed teams, spreadsheet work, campaign and content production, and any business that wants one subscription to cover most requests.                                                                                                                                            |
| Watch out for     | Personal plans train on your data by default. The Business tier does not. Put the team on Business.                                                                                                                                                                             |

## [ 2.4 Perplexity

Perplexity is a different kind of product. Rather than owning a model, it runs your question through live web search and then has a frontier model write the answer, with numbered citations beside every claim. That one design decision makes it the only tool of the group you can audit quickly, which is exactly what you want for market sizing, competitor checks, supplier due diligence and regulatory questions.

| Perplexity        |                                                                                                                                                                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Modes             | Search for quick cited answers. Pro Search for messy questions. Research for a multi-page cited report. Labs, on higher tiers, to produce spreadsheets, dashboards and decks. Model Council, on Max, runs a query across several frontier models at once for high-stakes decisions. |
| Price             | Free tier with roughly five Pro searches a day. Pro US\$20 a month or US\$200 a year. Max US\$200 a month. Enterprise Pro US\$40 per seat a month, Enterprise Max US\$325. The Comet browser is free on every tier.                                                                 |
| Standout features | Inline citations on everything. Spaces, which are project workspaces where you can query your own uploaded files alongside the web. The Comet browser.                                                                                                                              |
| Best for          | Any question where you need to show your working, or where being six months out of date would be expensive.                                                                                                                                                                         |
| Watch out for     | It researches, it does not operate. It will not draft in your voice as well as Claude, and it cannot act inside your CRM or helpdesk.                                                                                                                                               |

## [ 2.5 Microsoft Copilot

Copilot comes in two forms and confusing them is the most common mistake business owners make. The free consumer app at copilot.com is a capable general assistant, and its conversations can be used for training by default. Microsoft 365 Copilot is the business product. It sits inside Word, Excel, PowerPoint, Outlook and Teams, and it answers using your organisation's own email, files and meetings while respecting the permissions each person already has.

That grounding is the whole argument for it. Asking "summarise everything we agreed with this client over the last six months" is trivial for Copilot and impossible for the others without you manually supplying the material. One correction worth carrying: the standalone consumer Copilot Pro subscription at US\$20 a month no longer exists for new customers. Microsoft folded those features into Microsoft 365 Premium. If a guide or a search result quotes you Copilot Pro, it is working from old information.

| Microsoft 365 Copilot |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| What you get          | Copilot inside each Office app, Copilot Chat across your tenant, meeting summaries in Teams, and Copilot Studio for building simple internal agents.                                                                                                                                                                                                                                                                                                                                                                        |
| Price                 | Copilot Chat, which is web-grounded only, is included with eligible Microsoft 365 subscriptions. Copilot Business, which adds grounding in your own company data, is US\$21 per user a month on annual billing, roughly US\$25 monthly, on top of a qualifying base plan. Promotional pricing near US\$18 has been offered more than once, so ask. Since 1 July 2026 there are bundled SKUs: Business Standard with Copilot at US\$23.50 and Business Premium with Copilot at US\$32. The enterprise add-on remains US\$30. |
| Best for              | Teams already standardised on Microsoft 365 who want AI working on their own documents rather than in a separate tab.                                                                                                                                                                                                                                                                                                                                                                                                       |
| Watch out for         | Two things. The real cost is the base plan plus the add-on, not the add-on alone. And Copilot will surface anything a user already has permission to open, so audit your SharePoint and OneDrive sharing before you switch it on for everyone.                                                                                                                                                                                                                                                                              |

## [ 2.6 Gemini, from Google

Gemini is Google's assistant, and like Copilot it comes in two guises that matter to a business. The free consumer app is a strong general assistant whose conversations can be used to improve Google's models by default. Gemini in Google Workspace is the business product: it sits in the side panel of Gmail, Docs, Sheets, Slides and Meet and works on your own organisation's content, with your data contractually excluded from training. If your company already lives in Google Workspace, this is the one that fits how you already work, exactly as Copilot does for a Microsoft shop.

| Gemini            |                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Model tiers       | Flash for fast, high-volume work. Pro as the everyday default. A top reasoning mode, Deep Think, on the Ultra tier. Deep Research goes away and returns a cited report. Version numbers move, the Gemini 3 family in mid-2026, but the tier logic is stable.                                                                                                                                   |
| Price             | Free tier runs the Flash-class model. Google AI Pro US\$19.99 a month adds the Pro model, a one-million-token context, Deep Research, and image and video generation. Google AI Ultra US\$99.99 a month unlocks the top reasoning model and the highest limits. For business, Gemini is now bundled into Google Workspace Business and Enterprise plans rather than sold as a separate add-on. |
| Standout features | Gems, which are saved custom assistants, like a reusable brief. Deep Research. Canvas for side-by-side drafting. The side panel inside Gmail, Docs, Sheets, Slides and Meet. NotebookLM for querying your own document sets. Native image and video generation, and grounding in Google Search.                                                                                                |
| Best for          | Businesses already on Google Workspace who want AI inside Gmail, Docs and Drive working on their own content; strong general assistance and research for everyone else.                                                                                                                                                                                                                        |
| Watch out for     | The consumer app uses your chats to improve Google's models, with human review, unless you turn Gemini Apps Activity off. Keep company work inside Workspace, and do not connect a personal Gemini account to company data.                                                                                                                                                                    |

## [ 2.7 Side by side

Capability first, as a quick read; the commercial facts follow underneath.

|                        | Claude  | ChatGPT | Perplexity | Copilot | Gemini  |
|------------------------|---------|---------|------------|---------|---------|
| Long-form writing      | ● ● ● ● | ● ● ● ○ | ● ● ○ ○    | ● ● ○ ○ | ● ● ● ○ |
| Long-document analysis | ● ● ● ● | ● ● ● ○ | ● ○ ○ ○    | ● ● ● ○ | ● ● ● ○ |
| Cited research         | ● ● ● ○ | ● ● ● ○ | ● ● ● ●    | ● ● ○ ○ | ● ● ● ○ |
| Spreadsheet analysis   | ● ● ○ ○ | ● ● ● ● | ● ○ ○ ○    | ● ● ● ○ | ● ● ● ○ |

*More filled dots = stronger for that job. Empty = weak or not offered.*

|                               | Claude         | ChatGPT        | Perplexity     | Copilot        | Gemini              |
|-------------------------------|----------------|----------------|----------------|----------------|---------------------|
| Works on your company data    | Via connectors | Via connectors | Uploaded files | Natively, M365 | Natively, Workspace |
| Image generation              | No             | Yes            | No             | Yes            | Yes                 |
| Typical individual price      | US\$20         | US\$20         | US\$20         | US\$21 + base  | US\$19.99           |
| Trains on data, free tier     | Yes, opt out   | Yes, opt out   | Check settings | Yes, opt out   | Yes, opt out        |
| Trains on data, business tier | No             | No             | No             | No             | No                  |

Figure 1. Capability ratings are EncubIQ's editorial read for typical SMB use; commercial facts are as at 19 July 2026, quoted in US dollars at list rates for monthly billing. What you actually pay will differ: vendors bill in local currency, annual commitments run cheaper, non-profit and education rates exist, and sales tax or GST is added on top and is recoverable if you are registered. Canadian buyers should expect CAD billing. Confirm current figures on the vendor pricing pages before you buy anything.

## [ 2.8 Three sensible starting stacks

| Business shape                                       | Recommended stack                                                                                                                                      | Roughly per month                         |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| Solo operator or founder-led consultancy             | One paid workhorse, either Claude Pro or ChatGPT Plus, plus Perplexity Pro. Turn off training on both.                                                 | US\$40                                    |
| Two to ten people, not on Microsoft 365 or Workspace | ChatGPT Business for everyone, Perplexity Pro for whoever does research and bids. Add Claude Pro for the person who writes most.                       | US\$25 per seat plus extras               |
| Already on Microsoft 365                             | Microsoft 365 Copilot Business for roles that live in Office and Teams. Perplexity for research. Start with the twenty per cent who will use it daily. | Base plan plus ~US\$18 to US\$21 per seat |
| Already on Google Workspace                          | Gemini for Workspace for roles that live in Gmail and Docs. Perplexity for research. Same rule: licence the daily users first.                         | Included in Workspace plan tiers          |

### BUY NARROW, THEN WIDEN

The most common waste we see is buying seats for the whole team in month one. Licence the handful of people who will use it every day, prove the time saved, and expand from evidence rather than enthusiasm.

## 3 Choosing the right model for the job

Every vendor now sells the same ladder under different names: something fast and cheap, something balanced, something slow that reasons through problems step by step, and a research mode that goes away for several minutes and comes back with a report. Learning the ladder is more durable than learning this month's model names.




Figure 2. The escalation ladder. Start at the everyday tier and move up only when the answer is not good enough.

The tiers carry different names at each vendor, but they map onto each other cleanly.

| Tier      | Claude        | ChatGPT            | Gemini        |
|-----------|---------------|--------------------|---------------|
| Fast      | Haiku         | fast default       | Flash         |
| Everyday  | Sonnet        | the everyday model | Pro           |
| Reasoning | Opus          | Thinking           | Deep Think    |
| Research  | Research mode | Deep Research      | Deep Research |

Perplexity and Copilot ride on these same models rather than shipping their own tiers, so the logic still applies when you use them.

### [ 3.1 The overthinking tax

There is a real cost to reaching for the most powerful setting by default. Reasoning models are slower, they burn through usage limits faster, and on simple questions they produce longer, more hedged answers than you wanted. Asking a reasoning model for your opening hours rewritten for a Facebook post is like booking a structural engineer to hang a picture.

The inverse mistake is more expensive. Using a fast model to work through a pricing decision or a cash flow question produces something that reads confidently and is wrong. That is the situation worth paying the extra ninety seconds to avoid.

| If the task is                                         | Use                                  | Because                                               |
|--------------------------------------------------------|--------------------------------------|-------------------------------------------------------|
| Rewriting, shortening, tone changes, formatting        | Fast or everyday tier                | There is nothing to reason about                      |
| Drafting an email, post, description or job ad         | Everyday tier                        | The work is comprehension, not reasoning              |
| Summarising a long document or a meeting               | Everyday tier with the file attached | Quality of writing matters more than depth of thought |
| A pricing change, a hiring tradeoff, a cash flow model | Reasoning tier                       | Multi-step logic is where the fast models break       |
| Checking a fact, a competitor, a regulation            | Perplexity, or research mode         | You need citations you can click                      |
| Market scan, supplier landscape, category review       | Research mode                        | Depth across many sources beats a single answer       |

#### A HABIT WORTH FORMING

When an answer disappoints you, change one thing at a time. First improve the prompt, because that fixes it about eighty per cent of the time. Only then escalate the model. Most people do this the other way round and conclude the tool is weak when the brief was.

## 4 The anatomy of a prompt that works



A good prompt is a brief. If you would not hand the same instruction to a competent freelancer and expect usable work back, do not expect it here either. Six elements cover almost every case, and you will not need all six every time.

|             |                                                                              |
|-------------|------------------------------------------------------------------------------|
| ROLE        | "You are a service manager at an independent RV dealership."                 |
| TASK        | "Write a reply to this customer complaint."                                  |
| CONTEXT     | "They have spent \$40k with us over four years. The repair was late."        |
| CONSTRAINTS | "Under 150 words. No apology for anything we did not do. No discount offer." |
| FORMAT      | "Plain email. Subject line, three short paragraphs, sign-off."               |
| EXAMPLES    | "Match the tone of the two replies I have pasted below."                     |

Figure 3. Six components of a working prompt. Role, task and context do most of the work.

## [ 4.1 A simpler version to remember

Microsoft teaches a four-part formula in its Copilot training that is easier to hold in your head: goal, context, source, expectations. Google teaches almost the same thing for Gemini: persona, task, context, format. What you want, why and for whom, which file or email or meeting to use, and what the output should look like. If six elements feels like too much on a Tuesday morning, use four. The source element in particular is the one people skip and the one that most improves the answer.

## [ 4.2 Say what you do not want

Negative constraints are underused and unusually effective. These models have a default register: enthusiastic, slightly breathless, fond of the phrase "in today's fast-paced world". You can simply forbid it. A line such as "no marketing cliches, no exclamation marks, do not open with a rhetorical question, do not use the words leverage, unlock or seamless" changes the output more than any amount of praise for the brand.

## [ 4.3 Put long material first, ask last

When you are pasting in a long document, put the document at the top of the message and your question at the bottom. Anthropic's own guidance notes that with long inputs this ordering measurably improves the response. It costs nothing to do and it is easy to get backwards.

### THE THREE-LINE UPGRADE

If you take one thing from this section: add who you are, who it is for, and what it must not do.

Three extra lines, and the difference in output is larger than switching to a more expensive model.

## 5 Eight prompts, rewritten

These are drawn from work small businesses actually put in front of these tools. In each case the weak version is not a straw man. It is what most people type.

### [ 5.1 A social post for a promotion

The most common request in any small business, and the one that most reliably produces something you would never post.

#### WEAK PROMPT

Write a LinkedIn post about our summer sale.

*What comes back:* a post opening with "Exciting news!", three emoji, a bulleted list of benefits nobody asked for, and four hashtags.

#### STRONGER PROMPT

You write for a family-run outdoor gear shop in Squamish. Our customers are weekend hikers aged 30 to 55, mostly local, and they dislike being sold to.

Write a LinkedIn post announcing 20% off last season's packs from 1 to 14 August.

Rules: under 120 words. No emoji, no hashtags, no exclamation marks. Do not open with a question. Lead with why we are discounting, which is that we are clearing space for the winter range. Plain, dry, slightly self-deprecating.

Give me three versions with different opening lines.

Why it works. It supplies an audience, a reason to believe, a banned list, and asks for options rather than one answer. Asking for three versions costs nothing and roughly triples the chance one of them is usable.

### [ 5.2 A reply to an unhappy customer

Here the risk is not a bad sentence. It is accidentally admitting liability or offering money you did not intend to.

#### WEAK PROMPT

Write a nice reply to this angry customer email.

*What comes back:* a grovelling apology that accepts fault, offers a full refund, and promises a callback from a manager you have not spoken to.

STRONGER PROMPT

You are the owner of a 12-person landscaping company. Below is a complaint email from a residential client.

[paste the email]

Draft my reply. Constraints:

- Acknowledge the delay, which was real, without accepting fault for the drainage issue, which was pre-existing.
- Do not offer a refund, a discount, or free work.
- Offer one specific next step: a site visit on Thursday or Friday.
- Under 140 words. Warm but not apologetic in tone.

Then list separately anything in my draft that could be read as an admission of liability.

Why it works. The constraints are commercial, not stylistic. The final instruction turns the assistant into a reviewer of its own output, which catches the exact failure mode you are worried about.

## [ 5.3 A proposal or scope of work

Long, repetitive, high-stakes and largely templated. A good fit, provided you supply the raw material.

WEAK PROMPT

Write a proposal for a marketing project.

*What comes back:* a generic six-page document full of square brackets, invented deliverables, and a timeline unrelated to how you work.

STRONGER PROMPT

Attached: my standard scope of work template, notes from the discovery call, and two previous proposals that won.

Draft a scope of work for this prospect using my template structure exactly. Pull the deliverables and phasing from the call notes only. Do not invent deliverables I did not discuss.

Match the tone of the two winning proposals.

Where you do not have enough information, insert [NEEDS INPUT] and list those gaps at the end rather than guessing.

Why it works. Three instructions do the heavy lifting: use my structure, use only my source material, and flag gaps instead of filling them. That last one is the difference between a draft you edit and a draft you have to fact-check against yourself.

## [ 5.4 A job posting

Fast to produce and easy to get subtly wrong, since the default output is written for a company ten times your size.

### WEAK PROMPT

Write a job description for an office administrator.

*What comes back:* a corporate posting demanding five years of experience, a degree, and "the ability to thrive in a fast-paced environment".

### STRONGER PROMPT

We are a 9-person plumbing contractor in Surrey, BC. Hiring one office administrator, in person, four days a week, \$52k to \$58k.

The honest picture: the role is 60% scheduling and dispatch, 25% invoicing in QuickBooks, 15% fielding calls from frustrated customers. It suits someone calm who likes order. It does not suit someone who wants a career path into marketing.

Write the posting. Under 400 words. State the salary. No buzzwords, no "rockstar", no "fast-paced". Do not list more than five requirements. Write it so the wrong candidate self-selects out.

Why it works. Giving it the unglamorous truth about the role is what makes the posting distinctive. The instruction to help the wrong candidate self-select out is a strategic constraint, not a stylistic one, and it changes the whole document.

## [ 5.5 Making sense of a sales spreadsheet

Upload the file. Do not describe it, and do not paste numbers into the chat window and hope.

### WEAK PROMPT

Analyse my sales data and give me insights.

*What comes back:* observations you already knew, restated as findings. Revenue went up in December.

STRONGER PROMPT

Attached is 24 months of invoice-level sales data. Columns are date, customer, product line, units, revenue, and rep.

Answer four questions, showing the calculation for each:

1. Which product line has the worst revenue trend over the last 6 months compared with the same 6 months last year?
2. What share of revenue comes from our top 10 customers, and how has that concentration changed year on year?
3. Which customers bought in 2024 but not at all in 2025?
4. Is there a month-of-year pattern strong enough to plan around, or is what looks seasonal just two large orders?

Flag any data quality problems you hit. Do not round away detail that changes the conclusion.

Why it works. Specific questions beat "give me insights" every time. Asking it to show the calculation means you can check the arithmetic, and the fourth question explicitly invites it to talk you out of a pattern rather than confirm one.

## [ 5.6 Checking out a competitor

This is a Perplexity job, or a research mode job. It is the wrong task for a model working from memory.

WEAK PROMPT

Tell me about my competitors in the RV dealership market.

*What comes back:* a plausible summary that may be two years out of date, with no way to tell which parts are current.

STRONGER PROMPT

Research the four RV dealerships operating in the Fraser Valley listed below. For each, report: brands carried, whether they offer service and parts, published financing partners, current promotions, Google review count and average rating, and any ownership or acquisition news since January 2025.

Use only sources published in the last 18 months. Cite every claim. Where you cannot find something, say so rather than inferring it. Present as a comparison table.

Why it works. A date boundary, a named list of fields, a citation requirement, and explicit permission to return nothing. That last instruction is what stops a research tool quietly filling holes with reasonable-sounding guesses.

## [ 5.7 Turning meeting notes into decisions

A summary is not what you need. You need the decisions, the owners, and the things that were raised and dropped.

WEAK PROMPT

Summarise this meeting transcript.

*What comes back:* a chronological retelling of the meeting, which is no faster to read than the meeting was to attend.

STRONGER PROMPT

From the attached transcript of Tuesday's ops meeting, produce four sections:

1. Decisions made. One line each, with who decided.
2. Actions. Owner, action, and the date if one was stated. Write "no date agreed" where none was.
3. Open questions that were raised and not resolved.
4. Anything where two people appeared to disagree and the disagreement was not settled.

Do not include general discussion. If a section is empty, say so.

Why it works. The fourth section is the valuable one and nobody asks for it. Unresolved disagreement is what causes the same meeting to happen again in three weeks.

## [ 5.8 Writing an SOP from what is in your head

The highest-leverage use of these tools for most owner-managed businesses, and the least glamorous.

WEAK PROMPT

Write a standard operating procedure for onboarding new clients.

*What comes back:* a generic ISO-flavoured document describing a company that is not yours.

STRONGER PROMPT

I am going to describe how we onboard a new client. It will be disorganised because I am doing it from memory.

[dictate or type it however it comes out]

Turn this into a written SOP a new hire could follow without asking me questions. Number the steps. Note who does each step.

Then, separately: list the points where my description was vague or where I skipped what happens if something goes wrong. Ask me up to eight questions to close those gaps, and wait for my answers before revising.

Why it works. It takes the pressure off you to be organised, and it turns the assistant into an interviewer. Being asked eight good questions about your own process is frequently more useful than the document

that results.

## 6 Five habits that compound

Prompt wording matters less than most guides suggest. Working habits matter more. These five change the quality of everything you get back, and none of them takes more than a minute to adopt.

### [ 6.1 Give it the source material

The single most reliable way to stop an assistant inventing things is to give it the facts and tell it to stay inside them. Upload the contract rather than describing the contract. Paste the email thread. Attach the spreadsheet. Then add the sentence that does the real work: "use only the information provided; if something is not in the material, say so."

### [ 6.2 Show it two examples

If you want output in your voice, paste in two things you have written that you were happy with and ask it to match them. Two good examples beat two paragraphs of adjectives describing your brand voice. This is the technique every vendor calls few-shot prompting, and it is the fastest route out of generic output.

### [ 6.3 Make it critique its own work

Once you have a draft, the highest-value follow-up in this entire guide is some version of: "review that for weak claims, missing information and anything a sceptical reader would push back on, then rewrite it." The second pass is consistently better than the first, and it costs one line. Ask specifically for what worries you, whether that is legal exposure, factual accuracy or a client who will read it uncharitably.

### [ 6.4 Stop trying to write the perfect prompt

Treat the first output as a starting position. "Too long, cut it by half." "The second paragraph is the strongest, build around it." "This reads like a press release, make it sound like an email from a person." Three rounds of blunt feedback gets you further than any amount of upfront prompt engineering, and it is how people who are good at this actually work.


### [ 6.5 Set the context once

Every one of these tools now has somewhere to store standing context: Projects in Claude and ChatGPT, Spaces in Perplexity, Gems in Gemini, and custom instructions and memory in all of them. Write two hundred words describing your business, your customers, your tone and your non-negotiables, and put it there. You will stop re-explaining yourself and the baseline quality of every future answer rises.

#### IF YOU DO NOTHING ELSE

Adopt habits one and three. Attach the source material, then ask it to criticise its own draft. Those two alone remove most of what people dislike about AI output.

## 7 Where it goes wrong



These tools fail in characteristic ways. Once you can name the failure modes you will spot them in seconds, which is what makes it safe to use them for real work.

| Failure mode         | What it looks like                                                                                                         | What to do                                                                                 |
|----------------------|----------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Fabrication          | A statistic, a date, a case name or a quotation that sounds exactly right and does not exist. Confident, fluent, invented. | Require sources. Verify anything you would be embarrassed to be wrong about.               |
| Fabricated citations | A reference with a real-looking author, journal and year that leads nowhere when you search for it.                        | Click every link. If there is no link, treat the citation as unverified.                   |
| Agreeing with you    | You push back on a correct answer and it folds immediately, then agrees with your wrong version.                           | Ask "what is the strongest argument against what I just said?" rather than "are you sure?" |
| Confident wrongness  | No hedging on the parts it is least sure about. The tone is identical whether it knows or is guessing.                     | Ask it to rate its own confidence and mark which claims it cannot support.                 |
| House style          | Bland, symmetrical, faintly promotional prose. Lists of three. Words like leverage, unlock and seamless.                   | Ban the vocabulary explicitly and give it two samples of your own writing.                 |

| Failure mode      | What it looks like                                                                              | What to do                                                               |
|-------------------|-------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Quiet gap-filling | It does not have a number so it produces a plausible one rather than telling you it is missing. | Instruct it to insert a placeholder and list gaps instead of estimating. |

## [ 7.1 The cautionary tale worth knowing

Courts in several countries have now sanctioned lawyers who filed briefs citing cases that did not exist. The citations had real-sounding names, plausible reporter numbers and confident summaries. Nobody checked. It is worth knowing this happened in a profession whose entire training is in verifying sources, because the lesson is not that lawyers were careless. It is that fabricated material is convincing precisely when you are busy and it is telling you what you expected to hear.

## [ 7.2 A verification rule you will actually follow

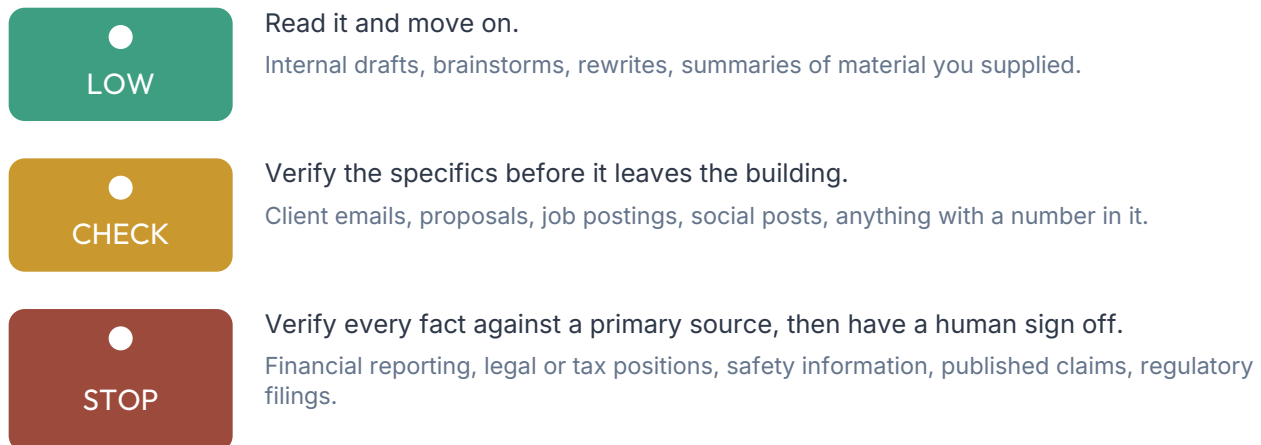


Figure 4. Three levels of verification. The middle band is where most business output sits.

### THE ACCOUNTABILITY POSITION

No vendor takes responsibility for what you send to a customer. If it goes out under your name, it is yours. Build one human check into any workflow you automate, and put it at the point of no return rather than at the start.

## [ 7.3 Preventing invention in the first place

Catching a fabrication after the fact is the expensive way to handle this. Most of it can be stopped before it reaches you, and none of the following costs anything.

Turn web search on for anything factual. A model answering from memory alone is reconstructing what a plausible answer looks like. The same model with search enabled goes and reads. If your question involves a price, a date, a regulation, a company or anything that has changed recently, turn search on before you ask, not after you doubt the answer.

Ask where to find a number, not for the number. Requests for statistics are the highest-risk thing you can type. "What is the average customer acquisition cost in my industry" invites a confident invention. "Which published sources report customer acquisition cost benchmarks for independent RV dealers, and what does each one actually say" gets you somewhere you can check.

Route factual questions to the tool that cites. If the answer needs to be defensible, ask Perplexity rather than a chatbot, then bring the sourced answer back to your writing tool. This is the practical reason for running two subscriptions rather than one, and it costs about twenty dollars a month to remove a whole category of risk.

Give it permission to fail. These systems default to producing an answer. Say the words: "if you do not know, say so; if the material does not contain it, write UNKNOWN rather than estimating." It measurably reduces gap-filling, because you have made admitting ignorance an acceptable output rather than a failure.

Cross-check anything consequential against a second tool. Two different systems rarely invent the same false fact. Paste the claim into the other one and ask it to verify. Thirty seconds, and it catches most of what would embarrass you.

Start a fresh chat for factual work. Quality degrades in long threads. Earlier mistakes stay in the context and get treated as established fact. If accuracy matters, open a new conversation and supply the material again.

#### THE ASSUMPTION WORTH DISCARDING

A slower, more expensive reasoning model does not hallucinate less about obscure facts. It reasons better and it sounds more authoritative, which makes its inventions harder to spot, not easier. Grounding beats model choice every time.

## [ 7.4 The pre-flight check

One page. Run it before anything leaves your business with your name on it. Most of it takes under a minute.

| Before you ask                                                                          | Why it matters                                           |
|-----------------------------------------------------------------------------------------|----------------------------------------------------------|
| 1 Turn web search on if the answer depends on facts, prices, dates or anything current. | Stops the model reconstructing from memory.              |
| 2 Attach the source material rather than describing it.                                 | Removes the need to invent the details you left out.     |
| 3 Add: use only what I have given you; write UNKNOWN for anything missing.              | Makes admitting a gap an acceptable answer.              |
| 4 Ask for sources alongside any claim you would repeat to a customer.                   | Unsourced claims are the ones that turn out to be false. |

| Before you send                                                                           | Why it matters                                     |
|-------------------------------------------------------------------------------------------|----------------------------------------------------|
| 5 Click every citation. No link means unverified, not true.                               | Fabricated references look exactly like real ones. |
| 6 Check every number, date, name and quotation against a primary source.                  | These are the specific things models invent most.  |
| 7 Ask it which claims it cannot support, then look at those first.                        | Cheaper than finding out from a customer.          |
| 8 Read it once as the recipient. Does it say anything you cannot personally stand behind? | You are accountable for it either way.             |

Figure 5. The pre-flight check. Steps one to four prevent, five to eight catch.

**IF YOU ONLY REMEMBER ONE LINE**

Anything you would be embarrassed to be wrong about in front of a customer, a regulator or your accountant gets verified against a source you can open. Everything else is a draft.

## [ 7.5 The bias problem, briefly

Hallucination is the failure everyone talks about. Bias is the quieter one, and for a business it carries more legal exposure. These systems learned from published text, which means they inherited its default assumptions. Three show up in ordinary business use.

Default demographics in marketing copy. Ask for a customer persona and you will often get a version of the same person: professional, urban, middle income, culturally Western. Ask for imagery and the pattern is sharper still. If your actual customer base does not look like that, the copy will land flat and you will not immediately know why. The fix is specificity. Describe your real customers rather than asking it to imagine some.

Screening and hiring. This is the one with genuine legal weight. Using these tools to rank, score or filter applicants imports assumptions you cannot inspect and cannot defend if challenged. Ontario now requires employers to disclose the use of AI in screening for publicly advertised roles, and it will not be the last jurisdiction to act. Use them to write the posting and to prepare interview questions. Do not use them to decide who gets interviewed.

Geographic and regulatory default settings. Unless told otherwise, these tools answer as though you are in the United States. For a Canadian business that means American employment norms, American tax treatment, American spelling and American consumer law, delivered with complete confidence. Put your jurisdiction in your standing context and it largely goes away.

#### THE PRACTICAL TEST

Before anything customer-facing goes out, ask one question: would a customer who is not the default assumption recognise themselves in this? If the honest answer is no, you have found the bias without needing a framework to name it.

## 8 Twelve jobs worth trying this month

Pick three. Preferably three you do every week, because repetition is where the time saving actually shows up. A task you do twice a year is not worth building a habit around.

| Function         | The job                                                                                   | Best tool                 |
|------------------|-------------------------------------------------------------------------------------------|---------------------------|
| Sales            | Turn discovery call notes into a first-draft proposal using your own template             | Claude                    |
| Sales            | Rewrite one outreach email for six different prospect types without it sounding automated | ChatGPT or Claude         |
| Marketing        | Produce a month of social posts from one long piece you already published                 | Any                       |
| Marketing        | Rewrite your service page for the customer who is comparing you against two others        | ChatGPT or Claude         |
| Customer service | Draft replies to your ten most common enquiries, then keep them as templates              | Any                       |
| Customer service | Turn a year of support emails into a written FAQ and a refund policy                      | Claude, Copilot or Gemini |
| Finance          | Interrogate your sales export for customer concentration and lapsed accounts              | ChatGPT                   |
| Operations       | Explain a supplier contract's payment and termination terms in plain English              | Claude                    |
| Hiring           | Write the job posting, the screening questions, and a scorecard for the interview         | Any                       |
| Operations       | Convert what only you know into written SOPs for the five jobs you cannot delegate        | Claude                    |
| Research         | Build a competitor comparison table with citations you can check                          | Perplexity                |

| Function | The job                                                                        | Best tool         |
|----------|--------------------------------------------------------------------------------|-------------------|
| Meetings | Extract decisions, owners and unresolved disagreements from recurring meetings | Copilot or Gemini |

#### HOW TO TELL IF IT IS WORKING

Measure one thing: hours back per week, self-reported by whoever does the task. If nobody can name a task that got faster after sixty days, the problem is the workflow you chose, not the software.

## 9 The rules of the road



This is the section most guides skip and most business owners need. It is short, and none of it is difficult.

### [ 9.1 Does it train on your data?

The pattern across all five vendors is the same and it is the most important thing in this guide. On free and personal plans, your conversations are generally used to improve the models unless you turn that off. On business and enterprise plans, they are contractually excluded. That single distinction is why a paid business tier is not a luxury.

| Platform | Personal or free tier                                                                                                                                                                           | Business or enterprise tier                        |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Claude   | Since August 2025 conversations are used for training unless you opt out at Settings, Privacy, Help improve Claude. Retention extends to five years if you allow it, thirty days if you do not. | Excluded from training under the commercial terms. |

| Platform          | Personal or free tier                                                                                                                                                  | Business or enterprise tier                                                                            |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| ChatGPT           | Used to improve models by default. Opt out at Settings, Data Controls. Temporary Chat never trains.                                                                    | Business, Enterprise, Edu and API data is not used for training.                                       |
| Gemini            | Consumer app conversations are used to improve Google's models, with human review, unless you turn off Gemini Apps Activity.                                           | Google Workspace content is excluded from training under the enterprise terms.                         |
| Perplexity        | Governed by your account data settings. Check them.                                                                                                                    | Enterprise tiers exclude your files and queries by default.                                            |
| Microsoft Copilot | Consumer Copilot conversations can be used for training by default, with an opt-out and an 18-month retention default. The EEA and several other markets are excluded. | Microsoft 365 Copilot: your prompts, responses and Graph data are not used to train foundation models. |

Figure 6. Training policies as at July 2026. Vendors change these; re-check annually.

## [ 9.2 What never goes into a personal account

| SAFE ON ANY PAID BUSINESS TIER     | NEVER PASTE INTO A PERSONAL OR FREE ACCOUNT |
|------------------------------------|---------------------------------------------|
| Your own drafts and internal notes | Customer names with their data attached     |
| Published or public material       | Employee records and HR files               |
| Anonymised or aggregated figures   | Health information of any kind              |
| Your marketing and website copy    | Bank details, card numbers, credentials     |
| Documents you own outright         | Client documents held under NDA             |
| Job descriptions and internal SOPs | Trade secrets and unpublished IP            |

Figure 7. The trust boundary. When in doubt, redact the names and the numbers before you paste.

## [ 9.3 Telling people you used AI

Disclosure norms have tightened and they now vary by where your customers are. Three points a small business should know:

- European Union. The AI Act's transparency obligations apply from 2 August 2026. In broad terms, AI-generated content must be marked as such, deepfakes must be disclosed, and some AI-generated text published to inform the public must be labelled, with a carve-out where a human has meaningfully reviewed and taken editorial responsibility. Penalties under this article run to 15 million euros or three per cent of worldwide annual turnover, whichever is higher. This matters if you sell into the EU.

- United States. There is no AI-specific disclosure statute. The Federal Trade Commission has instead applied existing rules, treating undisclosed AI-generated endorsements and fake reviews as deceptive, and has brought actions on that basis.
- Canada. The proposed federal AI legislation died when Parliament was prorogued in January 2025, so governance runs through privacy law and voluntary codes. Provincially, Ontario now requires employers to disclose the use of AI in screening for publicly advertised jobs, in force since 1 January 2026.

The workable position for a small business is simpler than the regulation. Disclose where a reasonable customer would care: synthetic images of people, anything presented as a testimonial or review, and automated responses a customer might mistake for a person. Keep a named human editorially responsible for everything you publish. You do not need to announce that a draft email was written with assistance, any more than you announce the use of a spellchecker.

### [ 9.4 A one-page policy you can adopt today

Write these eight lines, put them in your shared drive, and have everyone read them. It is not sophisticated governance and it does not need to be.

- 1 These are the approved tools. Do not use others for company work.
- 2 Use your company account. Personal accounts are not permitted for company material.
- 3 Never enter customer personal data, employee records, health information, credentials, or anything covered by an NDA.
- 4 Anonymise before you analyse. Strip names and identifiers first.
- 5 You are responsible for what you send. Check facts, numbers, names and dates before anything leaves the business.
- 6 Never paste AI-generated citations, statistics or legal statements without verifying them at the source.
- 7 Tell a manager before using AI in anything customer-facing and automated, such as a chatbot or an auto-reply.
- 8 If you are unsure whether something is allowed, ask before you paste, not after.

## 10 Your first thirty days



A plan you can run alongside actual work. The order matters more than the pace: settle the tooling and the rules before you scale usage, because retrofitting governance onto an established habit is much harder than starting with it.

- 1 Week 1 to 2**  
Choose one workhorse and one research tool. Put everyone on the paid business tier. Turn off training on any personal accounts still in use. Circulate the eight-line policy from 9.4.  
*Done when: Nobody is doing company work in a free personal account.*
- 2 Week 3 to 4**  
Write your standing context once, two hundred words on the business, customers, tone and non-negotiables, and load it into Projects, Spaces, Gems or custom instructions for every user.  
*Done when: Nobody has to re-explain the business at the start of a chat.*
- 3 Month 2**  
Pick three recurring tasks from section 8. Build and save a good prompt for each, using the six components in section 4. Test them on real work, not examples.  
*Done when: Three saved prompts exist and produce output someone would actually send.*
- 4 Month 2, review**  
Review honestly. Which tasks got faster and by how much? Where did the output need so much editing it was not worth it? Cancel unused seats.  
*Done when: You can name, in hours, what you got back.*
- 5 Month 3 and beyond**  
Expand to the next three tasks. Connect the tool to your Drive, SharePoint or CRM if the tier permits it. Audit file permissions before enabling any data grounding.  
*Done when: Usage is spreading because people asked for it, not because you mandated it.*

#### THE BENCHMARK

If, after sixty days, nobody on your team can point to a specific weekly task that got materially faster, stop and diagnose. In almost every case the cause is the task chosen or the brief written, not the tool bought.

### [ 10.1 The honest summary

These tools will not transform your business. They will take a meaningful slice of the writing, summarising, reformatting and first-draft work off the desk of everyone who does it. For an owner that is often the difference between working on the business and working in it. For a team of twenty it is a few hours a head, which only becomes real if someone owns the standards. That is a smaller claim than you will read elsewhere, and it is one you can actually test in a fortnight.

### [ 10.2 What this guide will not fix

In the interest of being straight with you, here is what stays broken after you have done everything in these pages. None of it is a tooling problem, which is why no tool solves it.

- Tool sprawl. Six months in, three departments are on three different assistants, two of them personal accounts, and nobody can say what company material has left the building.
- Nobody owns the prompt library. The prompts that work live in one person's chat history. When that person leaves, output quality quietly returns to where it started and no one connects the two.
- No line back to money. You can feel that things are faster. You cannot show which tasks got cheaper, whether the seats are worth the spend, or what to cut. Time saved that nobody redeploys is not a saving.
- Permissions you never audited. The moment you point an assistant at your own files, it inherits whatever access your staff already had. Most businesses discover their sharing settings are wrong at exactly the wrong moment.
- The work that is actually expensive. Attribution, pricing, retention, the decisions that move the number. Better drafting does not touch any of it.

#### WHERE THE LINE SITS

Everything in this guide can be tested by one person in an afternoon. The five problems above cannot, because each one is a question of policy, ownership and measurement rather than technique. Knowing which of the two you are facing is most of the value of reading this.

## [ Appendix A: Templates to copy

Adapt the bracketed parts. These are deliberately plain; the value is in the structure, not the wording.

Standing context, paste once into Projects, Spaces, Gems or custom instructions

### TEMPLATE

About my business: [what you do, who for, where, how many people].  
Our customers: [who they are, what they care about, what they are comparing us against].  
My role: [title and what you personally do].  
Tone: [three concrete adjectives, plus what to avoid].  
Never do this: [claims you cannot make, words you hate, formatting you will not use].  
When you do not know something, say so rather than estimating.

General-purpose drafting brief

### TEMPLATE

You are [role] at [business]. You are writing for [audience], who currently [what they think or feel].

Task: [what you want].  
Source material: [attached, or pasted below].  
Constraints: [length, tone, banned words, what not to promise].  
Format: [structure you want back].

Use only the source material. Where information is missing, insert [NEEDS INPUT] and list the gaps at the end.

The review pass, run this on every draft that matters

### TEMPLATE

Review your draft against these questions, then rewrite it.

1. Which claims cannot be supported by the material I gave you?
2. What would a sceptical [client / regulator / candidate] challenge?
3. What is missing that the reader will expect?
4. Which sentences are filler and could be cut without loss?
5. Does anything read as an admission, a guarantee, or a promise I did not authorise?

List your findings first, then give me the revised version.

Research brief with citation discipline

TEMPLATE

Research [topic] as it stands in [month, year].

Report on: [field 1], [field 2], [field 3].

Only use sources published after [date]. Cite every claim.

Where you cannot find something, write "not found" rather than inferring it. Note where sources disagree.

Present as [table / short brief]. Flag anything that looks like marketing material rather than independent reporting.

## [ References and disclosures

### Sources

- US Chamber of Commerce and Teneo Research, Empowering Small Business: The Impact of Technology on US Small Business, August 2025. Survey of 3,870 small businesses conducted 6 to 26 June 2025.
- Salesforce, Small and Medium Business Trends Report, 6th edition, December 2024. IDC, The Business Opportunity of AI, commissioned by Microsoft, 2024.
- Protiviti and London School of Economics, Bridging the Generational AI Gap, October 2025. US Census Bureau, Business Trends and Outlook Survey, August 2025.
- Vendor documentation for pricing, features and data policies: Anthropic (Claude), OpenAI (ChatGPT), Perplexity, Microsoft 365 Copilot, and Google (Gemini Apps Privacy Hub and the Google Workspace Generative AI Privacy Hub).
- Google, prompting guidance for Gemini in Workspace; Microsoft, Copilot prompting guidance; Anthropic, prompt engineering guidance.
- EU AI Act, Articles 50 and 99 (transparency and penalties). FTC guidance on AI-generated endorsements. Ontario employer AI-disclosure rules, in force 1 January 2026.

### Disclosures

Model names and version numbers change monthly; the tier logic (fast, everyday, reasoning, research) is stable, but confirm specific versions and prices on each vendor's own page before you act. Adoption and ROI figures are drawn from self-reported surveys and are directional, not precise. Nothing here is legal advice; regulatory timelines are shifting, so confirm current obligations with counsel for your jurisdiction and industry. This edition is current as at 19 July 2026.

## [ About EncubIQ Consulting

EncubIQ Consulting is a management consulting practice working with small and mid-sized businesses on the unglamorous parts of growth: knowing which marketing actually worked, keeping customers longer, and putting technology to work without buying more of it than the business can absorb.

### Core practices

- Revenue Attribution and Digital Performance. Measurement that survives contact with reality, so you can tell which spend earned its keep and which did not.
- CRM and Customer Lifecycle Strategy. Retention, segmentation and lifecycle design for businesses that have outgrown a spreadsheet.
- AI Strategy and Integration. Practical adoption for owner-managed businesses, of the kind described in this guide, taken from policy through to working process.

The practice works across automotive and RV retail, hospitality and gaming, retail and e-commerce, professional services and consumer goods, serving clients in Canada, the United States and internationally.

### IF THIS WAS USEFUL

The natural next step after this guide is a short diagnostic: which three tasks in your business would repay this effort first, what your policy should say, and what you should not automate yet. That conversation takes about thirty minutes and does not require you to buy anything.

Simran Kohli, Founder and Principal  
EncubIQ Consulting  
[encubiq.com](https://encubiq.com) · [contact@encubiq.com](mailto:contact@encubiq.com)

© 2026 EncubIQ Inc. All rights reserved. EncubIQ and the [En]IQ device are trademarks of EncubIQ Inc. This guide may be shared in full and unaltered with attribution. Extracts may be quoted with credit to EncubIQ Consulting.