PART 0

Why You're Not Getting Results

(And What's About to Change)

The Professional's AI Playbook

The Practical Guide to Making LLMs Work for Your Career

THE 7 REAL PROBLEMS PROFESSIONALS FACE WITH AI CHATBOTS

This isn't theory. These are the actual frustrations reported by thousands of professionals across Reddit communities (r/ChatGPT, r/ClaudeAI, r/LocalLLaMA), McKinsey surveys, and Microsoft's Future of Work research. If any of them sound familiar, you're in the right place.

1. The Blank Page Syndrome

You open ChatGPT or Claude. The cursor blinks. You type something vague like "help me with my project." The response is generic, useless, or both. You close the tab and go back to doing things the old way.

The problem isn't willpower or intelligence. It's that nobody gave you a clear starting point tailored to your actual job.

2. Prompt Anxiety

You've heard that "prompt engineering" is the key to getting good results. But every time you try, you second-guess yourself. Should I give it a role? How much context? Am I being too vague? Too specific?

Most professionals type the bare minimum and get bare-minimum results. Then they conclude the tool isn't useful for their work.

3. Re-explaining Yourself Every Single Time

Every new conversation starts from zero. The AI doesn't know your job title, your industry, your clients, your writing style, or what you worked on yesterday. You spend more time setting context than getting value.

Imagine hiring a new assistant every morning who has amnesia. That's the default experience.

4. The Workflow Gap

You can see how AI helps with isolated tasks: rewriting an email, summarizing a document. But you can't figure out where it fits in your actual daily workflow. When exactly do you open it? What do you feed it? How does the output connect to your next step?

Knowing a tool exists and knowing how to integrate it into your work are two completely different skills.

5. Generic Tutorials, Specific Job

Every AI guide you've seen shows the same examples: write a poem, plan a birthday party, summarize an article. None of them show what an operations manager, a finance controller, or a freelance consultant actually needs.

The gap between "generic demo" and "my Tuesday morning" is where most professionals give up.

6. Fragmented Learning

There's a YouTube video here, a Twitter thread there, a Medium article somewhere else. Some say "use system prompts." Others say "build custom GPTs." A few mention "skills" or "connectors" without explaining what they are.

You end up with scattered knowledge, no coherent system, and the nagging feeling that you're missing something important.

7. No Quick Wins in Your Own Context

If you work in a company, nobody around you is sharing how they actually use AI for your type of work. If you're a freelancer, you're figuring it out alone. There's no colleague to ask "hey, how do you use this for client proposals?"

Without visible, relevant success stories, motivation fades fast.

> **Sound familiar?**
> If you recognized yourself in three or more of these, this guide was built for you. Every chapter directly addresses one or more of these problems with practical, immediately usable solutions.

THE 5 MISTAKES THAT MAKE 80% OF PROFESSIONALS GIVE UP

Before we fix anything, let's name what goes wrong. These aren't character flaws. They're design traps built into the way AI tools are presented today.

1. **Treating it like a search engine.** You type a question, read the answer, and leave. That's Google behavior. An LLM is a thinking partner, not a lookup tool. The value multiplies when you iterate, push back, and refine.
2. **Starting with the hardest task.** You try to make the AI write your entire business plan or redesign your financial model on day one. It fails (because it lacks context), and you write off the technology. Start with a 5-minute task you do every week.
3. **Never saving your context.** You had one great conversation where the AI perfectly understood your situation. Then you started a new chat and lost everything. Without a persistent context system, every session is a cold start.
4. **Copying prompts from the internet.** You paste a "viral prompt" from Twitter. It produces impressive-looking output that has nothing to do with your actual needs. Prompts aren't spells. They only work when adapted to your specific context, constraints, and goals.
5. **Waiting for the perfect use case.** You keep saying "I'll start using it when I find the right project." That day never comes. The professionals who get results started with something small, imperfect, and real. They iterated from there.

> **The pattern**
> Every mistake above shares the same root cause: no system. You've been given a power tool with no instruction manual tailored to your job. This guide is that manual.

WHAT THIS GUIDE WILL CHANGE FOR YOU

Here's the honest promise. By the time you finish this guide and implement its tools, you will have moved from "I occasionally ask ChatGPT random questions" to "I have a personalized AI system that saves me 5+ hours every week."

Specifically, you will:

- Understand which AI tool fits your work style and budget (ChatGPT vs. Claude vs. Gemini) and stop second-guessing your choice.
- Know how to write prompts that produce professional-grade outputs on the first or second try, using a repeatable method you can teach to others.
- Have a persistent context system set up so the AI already knows your role, your clients, your vocabulary, and your standards before you type a single word.
- Own 10 ready-to-use playbooks for tasks you actually do at work: meeting prep, email drafting, document analysis, proposal writing, and more.
- Produce polished deliverables (presentations, formatted reports, structured documents) directly from the chat, without touching PowerPoint or Word manually.
- Have a fully architected setup: projects, skills, connectors, and reference files, all configured for your specific profession.

What you'll walk away with

This isn't a reading-only PDF. Embedded in this guide are three interactive AI skills you can install and use immediately:

SKILL	WHAT IT DOES
RCFC Prompt Builder	Transforms your rough prompt into a structured, professional-grade prompt through guided questions (Role, Context, Format, Constraints).
AI Onboarding	Interviews you about your work and generates 5 ready-to-use context files so the AI knows you from the first message.

Setup My Own LLM	Analyzes your profession and designs your complete AI architecture: projects, skills, connectors, and reference files.

HOW TO USE THIS GUIDE

You have two options. Pick whichever fits your style.

Option A: The Linear Path (recommended for first read)

Read from Part 1 to Part 7 in order. Each chapter builds on the previous one. By the end, you'll have a complete, working AI setup tailored to your profession.

Option B: The "I Need This Now" Path

Jump directly to the chapter that matches your most pressing need:

IF YOU NEED TO…	GO TO…
Pick the right AI tool	**Part 1**
Write better prompts immediately	**Part 2**
Stop re-explaining yourself	**Part 3**
Make a presentation with AI	**Part 4**
Get a ready-made workflow for a specific task	**Part 5**
Build your full AI environment	**Part 7**

> **One more thing**
> This guide is not about becoming a "prompt engineer." It's about becoming a professional who happens to have an AI system that makes them dramatically more effective. The technology is the means. Your expertise is the point.

PART 1

Understand the Machine, Choose Your Tool

The Professional's AI Playbook

HOW AN LLM "THINKS" — THREE METAPHORS THAT ARE ALL YOU NEED

You don't need to understand transformer architectures or attention mechanisms. You need three mental models that will change the way you interact with any AI chatbot.

Metaphor 1: The Brilliant Intern

Imagine you just hired the smartest intern on earth. They graduated top of their class, read every textbook in every field, and can write in 50 languages. But they started this morning. They don't know your company, your clients, your preferences, or what you worked on last week.

If you say "handle this," they'll produce something generic and impressive-looking that misses the point entirely. If you say "here's who we are, here's the client, here's what good looks like, now handle this," they'll deliver something remarkable.

Takeaway: the quality of the output is directly proportional to the quality of the briefing. Always.

Metaphor 2: The Mirror

An LLM reflects back the quality of your input. Vague question? Vague answer. Structured, specific, well-contextualized request? Structured, specific, well-contextualized response.

This is why "prompt engineering" isn't about tricks or magic words. It's about learning to communicate clearly what you actually want. Incidentally, that's a skill that improves your human communication too.

Takeaway: if the output is bad, the first place to look is your input.

Metaphor 3: The GPS

A GPS doesn't know where you want to go until you tell it. Once you do, it gives you the fastest route. If you take a wrong turn, it recalculates. It doesn't judge you, and it doesn't get tired.

An LLM works the same way. Give it a destination (your desired output), and it will find a path. If the first result is off, say so. It will adjust. The best users aren't the ones who write the perfect prompt on the first try — they're the ones who iterate fastest.

Takeaway: iteration is not a sign of failure. It's the actual method.

WHAT IT DOES WELL, WHAT IT DOES BADLY

Before you start using an LLM for work, you need to calibrate your expectations. Not everything belongs in a chatbot. Here's an honest breakdown.

TASK TYPE	CONFIDENCE	NOTES
Drafting emails and messages	**High**	Especially with tone/style context
Summarizing long documents	**High**	Best with copy-paste, not file links
Brainstorming and ideation	**High**	Excels when given constraints
Structuring messy thoughts	**High**	One of the highest-ROI uses
Rewriting and editing text	**High**	Give examples of your style
Explaining complex concepts	**High**	Great for learning new domains
Data analysis (tables, CSV)	**Medium-High**	Strong with small/mid datasets
Creating presentations	**Medium-High**	Requires iteration on design
Writing code	**Medium-High**	Best for scripting, automation
Research and fact-finding	**Medium**	Always verify critical facts
Strategic decision-making	**Medium**	Use as sparring partner, not oracle
Precise calculations	**Low-Medium**	Use code execution, not chat math
Real-time / live data	**Low**	Needs web search or connectors
Confidential / regulated work	**Caution**	Check your company's AI policy

> **The golden rule**

> Use AI for the first 80% of any task (the draft, the structure, the raw material), then apply your expertise for the last 20% (judgment, accuracy, context). Never ship raw AI output without review.

CHATGPT VS. CLAUDE VS. GEMINI: WHICH ONE IS RIGHT FOR YOU?

In this chapter, we'll focus on the three dominant LLM platforms: ChatGPT (OpenAI), Claude (Anthropic), and Gemini (Google). They are the tools you'll encounter most in professional settings, and each has a distinct philosophy and ecosystem.

But before we compare them, here's something important that most AI guides won't tell you: the market is converging. Three years ago, choosing the "right" model mattered enormously — the gap between the best and the rest was massive. Today, that gap has narrowed dramatically. All major models are converging toward the same level of intelligence. They all handle writing, analysis, coding, and reasoning at a professional level. The differences are real, but they're increasingly about ecosystem, workflow integration, and personal preference — not raw capability.

> **From the author: stop watching the leaderboards**
> Don't listen to the noise. Don't spend your time comparing benchmark scores or refreshing model leaderboards. As I'm writing this guide, Claude Opus 4.7 is leading the way on most benchmarks. But if you take the time to compare it side-by-side with GPT-5.5 or Gemini 3.1 Pro on your actual work tasks, you'll see a difference — but nothing dramatic for daily professional use. Pick one, get good at it, and invest your energy in building a system around it. That's where the real leverage is — not in chasing the latest model release.

The Comparison Matrix

	CHATGPT	CLAUDE	GEMINI
Company	OpenAI	Anthropic	Google
Price (Pro tier)	$20/mo	$20/mo	$20/mo
Free tier	Yes (GPT-5.5 Instant)	Yes (limited)	Yes (generous)
Long writing quality	★★★☆	★★★★	★★★☆
Reasoning depth	★★★★	★★★★	★★★☆
Following instructions	★★★☆	★★★★	★★★☆
Coding	★★★★	★★★★	★★★☆
Data / spreadsheets	★★★★	★★★☆	★★★★

Image generation	GPT Image (native)	No	Nano Banana (native)
Web search	Built-in	Built-in	Built-in (Google)
File handling	Strong	Strong	Strong
Presentation creation	Limited	Built-in (.pptx)	Google Slides
Ecosystem	Plugins, GPT Store	Projects, Skills	Google Workspace
Best for	All-rounder, images	Writing, analysis	Google users, data

The Recommendation by Profile

You're an analytical professional (finance, consulting, operations)

Start with Claude. It excels at long-form reasoning, structured analysis, following complex instructions, and producing polished documents. Its Projects feature lets you build persistent workspaces per client or workstream. If you need image generation too, add ChatGPT as a secondary tool.

You're a creative professional (marketing, content, design)

Start with ChatGPT. The combination of strong writing, native image generation (GPT Image), and the GPT Store for specialized tools makes it the most versatile creative platform. Claude is a strong alternative if you prioritize writing quality over image generation.

You live in the Google ecosystem

Start with Gemini. If your company runs on Google Workspace (Gmail, Drive, Slides, Sheets), Gemini's native integration is unmatched. It can pull from your Drive, draft in your Gmail, and create Google Slides directly. The AI quality is slightly behind Claude and ChatGPT on complex tasks, but the ecosystem advantage can outweigh that.

You're not sure yet

Start with one. Any one. The skills you learn with one LLM transfer directly to the others: prompt structure, iteration, context management. Pick the one that fits your budget and ecosystem, get good at it, and switch later if needed. The worst decision is no decision.

> **Pro tip**
> Many power users run two tools: one primary (for daily work) and one secondary (for specific tasks the primary handles less well). A common combo is Claude for writing/analysis + ChatGPT for images and quick lookups, or Gemini for Google Workspace + Claude for deep work.

THE 6 MODES OF INTERACTION: FROM BASIC CHAT TO FULL AUTOPILOT

Most people only use Mode 1. The professionals getting 10x results are operating in Modes 3 through 6. Here's the full spectrum, from simplest to most powerful.

Mode 1: Simple Chat

You open the chatbot, type a question, get an answer, leave. This is how 90% of people use AI today.

- **When it's fine:** quick lookups, one-off questions, casual brainstorming.
- **When it's not enough:** any recurring task, anything requiring your specific context.

Available in: ChatGPT, Claude, Gemini — all platforms.

Mode 2: System Instructions

You give the AI a set of standing instructions that apply to every conversation. Think of it as a permanent briefing note pinned to the top of every chat.

- **What you set:** your name, role, industry, preferred tone, output format, things to avoid.
- **Impact:** the AI stops being generic. It knows you're a finance consultant who prefers concise bullet points, not a college student who wants essays.

Where to find it: ChatGPT → Settings → Custom Instructions | Claude → Settings → General → Instructions for Claude | Gemini → Settings → Extensions & Preferences.

Mode 3: Projects & Persistent Workspaces

Instead of one generic chat, you create dedicated workspaces for different areas of your work. Each project has its own instructions, reference files, and conversation history.

- **Example:** a "Client Proposals" project that knows your pricing structure, your boilerplate text, and your past proposals. Every new proposal conversation starts with all that context already loaded.
- **Impact:** you stop re-explaining yourself. The AI becomes a specialist for each area of your work.

Where to find it: Claude → Projects | ChatGPT → Custom GPTs / Projects | Gemini → Gems.

Mode 4: Skills

This is where it gets powerful. A skill is a reusable capability that the AI can trigger on demand. Instead of re-writing the same complex prompt every time, you define it once as a skill, and the AI activates it whenever you need it.

What is a skill, concretely?

A skill is a set of instructions (typically a text file) that tells the AI exactly how to handle a specific type of task. It includes the role to adopt, the questions to ask, the format to produce, and the constraints to respect.

Examples of professional skills:

- Meeting Notes Structurer — paste raw notes, get decisions, action items, owners, and deadlines.
- Proposal Drafter — answer 5 questions, get a formatted client proposal.
- Email Tone Adapter — paste a draft, specify the relationship and goal, get a rewritten version.
- Weekly Report Generator — input bullet points, get a structured status report in your company's format.

How to create a skill:

In Claude, you add a file to your Project with a specific structure (name, description, trigger phrases, and instructions). In ChatGPT, you create a Custom GPT with similar instructions. In Gemini, you configure a Gem. The concept is the same across platforms; only the implementation differs.

> **This guide includes 3 ready-made skills**
> You'll find the RCFC Prompt Builder (Part 2), the AI Onboarding skill (Part 3), and the Setup My Own LLM skill (Part 7) ready to install. Each comes with full instructions for Claude, with adaptation notes for ChatGPT and Gemini.

Mode 5: Connectors (MCP)

Until now, the AI lived in a bubble. It could only work with what you typed or pasted into the chat. Connectors break that wall.

What are connectors?

Connectors (technically called MCP — Model Context Protocol) allow the AI to directly access your external tools: Google Drive, Gmail, Notion, Slack, Google Calendar, and more. Instead of copy-pasting data into the chat, the AI reads and writes to your actual tools.

What this unlocks:

- "Search my Google Drive for the Q3 report and summarize the key findings."
- "Check my calendar for tomorrow and draft a prep email for each meeting."
- "Find the latest thread in Slack about the product launch and give me the status."
- "Create a Notion page with the meeting notes from today's call."

How to set them up:

Claude: the most straightforward. Go to your chat and click the connectors menu (or go to Settings → Connectors). You'll see a list of available services: Google Drive, Gmail, Google Calendar, Notion, Slack, and more. Click "Connect," authorize, and you're done. Once connected, the AI can access them in any conversation. This is the smoothest connector experience across all three platforms.

ChatGPT: also supports MCP connectors, but calls them "Apps." Go to Settings → Apps to browse and connect available services (Canva, Spotify, Notion, and others). For more advanced integrations, ChatGPT offers a Developer Mode (Plus and above) where you can connect custom MCP servers — but this requires a remote HTTPS endpoint, which means a bit more technical setup. The pre-built apps in the directory are one-click to connect; custom MCP servers are for power users.

Gemini: connects natively to Google Workspace (Gmail, Drive, Calendar, Docs, Sheets) without any setup — if you're logged into your Google account, it just works. For third-party connectors beyond Google's ecosystem, the experience is more limited. Gemini Enterprise offers some third-party

connectors (Jira, Salesforce, Confluence), but the catalog is still small compared to Claude or ChatGPT. Google is investing heavily in MCP support, but as of now, connecting non-Google tools to Gemini requires either an enterprise plan or developer-level setup.

> **Start small**
> You don't need to connect everything on day one. Start with the one tool where you waste the most time searching for information (usually Google Drive or email). Add more connectors as you discover the value.

Mode 6: Agents & Automation

This is the frontier. Agents are AI workflows that run multiple steps autonomously: they search, analyze, write, and act, with minimal human intervention. Think of them as the difference between sending one text message and having a personal assistant who handles your entire morning routine.

Examples of what agents can do:

- Monitor your inbox every Sunday morning, extract the key insights from your industry newsletters, and create a weekly briefing in Notion.
- Scan a client's latest financial report, compare it to the previous quarter, and draft an analysis with visualizations.
- Receive a voice memo via Telegram, transcribe it, extract the action items, and add them to your task manager.

Agents typically require tools beyond the basic chat interface: workflow platforms like n8n, Make, or Zapier; agent frameworks like Claude Code or custom pipelines; and some technical setup. We'll cover the concepts in Part 6 (Leveling Up) so you know what's possible, even if you're not ready to build them today.

The Progression at a Glance

MODE	WHAT IT IS	EFFORT TO SET UP	VALUE MULTIPLIER

1. Chat	One-off questions	Zero	**1x**
2. Instructions	Persistent preferences	10 minutes	**2x**
3. Projects	Dedicated workspaces	30 min per project	**5x**
4. Skills	Reusable capabilities	1 hour per skill	**10x**
5. Connectors	Linked external tools	5 min per connector	**15x**
6. Agents	Autonomous workflows	Hours to days	**50x+**

> **Where this guide takes you**
> By the end of Part 7, you'll be operating comfortably at Mode 4-5 (Skills + Connectors), with a clear understanding of how to reach Mode 6 when you're ready. That alone puts you ahead of 95% of professionals using AI today.

PART 2

The Art of Professional Prompting

Including the RCFC Skill

THE RCFC METHOD: THE ONLY PROMPT FRAMEWORK YOU NEED

Forget the dozens of "prompt engineering" frameworks floating around the internet. Most of them are overly complex, academic, or designed for AI researchers. As a working professional, you need exactly one method. It's called RCFC.

LETTER	STANDS FOR	WHAT IT MEANS
R	**Role**	Who should the AI pretend to be? What expertise should it bring?
C	**Context**	What's the situation? Who's the audience? What do you already know?
F	**Format**	What should the output look like? Length, structure, tone, style?
C	**Constraints**	What should it avoid? What are the boundaries and rules?

Why RCFC Works

Every bad prompt is missing at least one of these four elements. When you say "write me an email," you're giving zero R, zero C, zero F, and zero C. The AI fills in every blank with generic defaults, and you get a generic result.

When you say "You are an experienced project manager (R). I need to inform my client that their deliverable will be delayed by two weeks due to a vendor dependency (C). Write a professional, empathetic email of 150-200 words (F). Don't blame the vendor directly and don't make promises about the new deadline until I confirm it (C)," you get something you can actually send.

The Before and After

Before (no RCFC):

```
Write me an email about a project delay.
```

Result: a generic, formal, lifeless email that sounds like it was written by a robot. You rewrite it from scratch.

After (with RCFC):

```
Role: You are a senior project manager with 15 years of
experience in IT consulting.

Context: I need to email my client (CFO of a mid-size
manufacturing company) that the ERP migration deliverable
will be delayed by 2 weeks. The cause is a vendor dependency
(the hosting provider missed their SLA). The client
relationship is strong but they are under pressure from
their board.

Format: Professional but warm email, 150-200 words.
Subject line included. Use short paragraphs.

Constraints: Don’t blame the vendor by name. Don’t commit
to a new deadline (I’ll confirm separately). Acknowledge
the impact on their timeline. End with a concrete next step
(a call this week to discuss mitigation).
```

Result: an email you can send with minor edits. Saved you 15 minutes and produced something better than what you would have written under time pressure.

> **The 80/20 rule of prompting**
> You don’t need to fill every RCFC element with a paragraph. Sometimes one sentence per letter is enough. The point isn’t to write a dissertation — it’s to make sure none of the four elements is completely missing. Even a quick “Role: marketing writer. Context: LinkedIn post about X. Format: under 200 words. Constraints: no hashtags” is infinitely better than “write a LinkedIn post.”

WHY 90% OF PROMPTS FAIL (AND WHY ITERATION BEATS PERFECTION)

Here's something the prompt engineering influencers won't tell you: even with a perfect RCFC prompt, the first output is rarely the final one. And that's not a problem — it's the method.

The Iteration Mindset

Professional AI users don't write one prompt and accept the result. They work in a loop:

1. **Prompt** — send your RCFC-structured request.
2. **Evaluate** — read the output critically. What's good? What's off?
3. **Refine** — tell the AI exactly what to change. "The tone is too formal. Make it more conversational while keeping the key points."
4. **Repeat** — keep going until you're satisfied. Two to three rounds is typical.

The Three Most Common Prompt Failures

Failure 1: No role = no perspective

Without a role, the AI defaults to "generic helpful assistant." It won't think like a CFO, write like a copywriter, or analyze like a consultant. Always assign a role, even a simple one.

Failure 2: No constraints = no boundaries

If you don't tell the AI what NOT to do, it will take the path of least resistance. It will use clichés, add unnecessary caveats, pad the response with filler, or include information you didn't ask for. Constraints are where the output goes from "decent" to "usable."

Failure 3: One-shot thinking

You type one prompt, get one response, and judge the tool based on that single interaction. That's like judging a new employee based on their first hour. The magic is in the conversation: push back,

redirect, add nuance. The AI improves dramatically when you engage with it rather than interrogate it.

> **Practical rule**
> If you're spending more than 5 minutes crafting the "perfect" prompt, stop. Send what you have and iterate. A good-enough prompt followed by two rounds of refinement will beat a perfect prompt every time.

BEYOND RCFC: PROMPTING TECHNIQUES WORTH KNOWING

RCFC is your daily driver. It covers 80% of professional use cases. But some situations call for a more specific approach. Here are six advanced techniques that experienced AI users reach for when RCFC alone isn't enough.

You don't need to memorize all of them. Scan the list, note the ones relevant to your work, and come back when you need them.

1. Chain-of-Thought (CoT)

What it is: Ask the AI to reason step by step before giving its final answer. Instead of jumping to a conclusion, it shows its work.

When to use it: complex analysis, math, logic problems, multi-step decisions, anything where the reasoning matters as much as the answer.

How to trigger it:

```
Think through this step by step before answering.
Show your reasoning at each stage.

[your question or problem]
```

Example:

```
I'm deciding between hiring a full-time employee
vs. a contractor for a 12-month project.
Think through this step by step:
- Compare total cost (salary + benefits + overhead
  vs. daily rate × working days)
- Consider flexibility, IP ownership, and ramp-up time
- Give me a recommendation with your reasoning
```

Why it works: LLMs produce significantly better answers on complex tasks when forced to reason explicitly rather than jump to a conclusion. It also lets you spot where the reasoning goes wrong.

2. Few-Shot Prompting

What it is: Give the AI 2-3 examples of the input/output pattern you want, then ask it to follow the same pattern for your actual request.

When to use it: when the output format is very specific and hard to describe in words, or when you want the AI to match a particular style or structure.

How to trigger it:

```
Here are examples of the format I want:

Input: [example input 1]
Output: [example output 1]

Input: [example input 2]
Output: [example output 2]

Now do the same for:
Input: [your actual input]
```

Example:

```
I need to write one-line summaries for meeting notes.
Here's the style I want:

Input: 45-min call with marketing about Q3 campaign
   launch. Agreed on Sept 15 date, Sarah owns creative.
Output: Q3 campaign launches Sept 15; Sarah leads
   creative. No blockers.

Input: Budget review with CFO. Need to cut 10%.
   Decided to reduce travel and defer hiring.
Output: 10% budget cut confirmed; travel + hiring
```

```
   deferred. CFO sign-off pending.

Now do the same for:
Input: [paste your meeting notes]
```

Why it works: showing is more powerful than telling. Two good examples communicate format, length, tone, and level of detail more precisely than a paragraph of instructions.

3. Persona Stacking

What it is: Ask the AI to evaluate a situation from multiple expert perspectives, either sequentially or as a simulated panel.

When to use it: strategic decisions, risk assessment, anything where you need to see a problem from different angles before committing.

How to trigger it:

```
Analyze the following situation from three
perspectives:

1. As a [expert type 1]: what do you see?
2. As a [expert type 2]: what concerns you?
3. As a [expert type 3]: what opportunity do
   you spot?

Then synthesize: what's the best path forward
considering all three viewpoints?

[describe your situation]
```

Why it works: it forces the AI to consider trade-offs you might miss when thinking from a single perspective. Particularly powerful for decisions with financial, operational, and people dimensions.

4. Constraint Tightening

What it is: Deliberately over-constrain the output to force the AI into a specific, high-quality lane. The tighter the box, the more creative and precise the result.

When to use it: when outputs are too generic, too long, or too "AI-sounding." Adding constraints is the fastest way to improve quality.

How to trigger it:

```
Write [deliverable] with these constraints:
- Maximum [X] words
- Must include [specific element]
- Must NOT include [specific thing to avoid]
- Written for [specific reader who will judge it]
- Tone: [very specific tone descriptor]
- If you have to choose between [A] and [B],
  always choose [A]
```

Why it works: constraints eliminate the AI's default behaviors (padding, hedging, being overly formal). The more specific your constraints, the less editing you do afterwards. Think of constraints as a filter that removes everything you'd delete manually.

5. Iterative Refinement (Prompt Chaining)

What it is: Break a complex task into sequential steps, where each prompt builds on the output of the previous one. Instead of one mega-prompt, you have a conversation.

When to use it: long documents, multi-section reports, complex analyses, anything where trying to do everything in one shot produces mediocre results.

How to trigger it:

```
Step 1: "Create an outline for [document]."
Step 2: "Good. Now write section 1 in detail."
Step 3: "Now section 2. Keep the same tone."
Step 4: "Review the full document for consistency
   and flag any gaps."
```

```
Step 5: "Tighten the executive summary to under
   100 words."
```

Why it works: each step gets the AI's full attention. A 5-step chain almost always produces better results than a single prompt asking for everything at once. It also gives you control points where you can redirect before the AI goes too far in the wrong direction.

6. Pre-Mortem / Red Team

What it is: Ask the AI to attack its own output, find weaknesses, or imagine why a plan might fail before you execute it.

When to use it: before sending a proposal, launching a project, making a major decision, or presenting to stakeholders. It's your built-in devil's advocate.

How to trigger it:

```
You just produced [the output above].
Now put on a different hat.

Imagine you are [a skeptical client / a hostile
board member / a competitor / a risk auditor].

- What are the 3 biggest weaknesses in this?
- What would you attack first?
- What's missing that would make you take this
  more seriously?

Be brutally honest. I need the criticism now,
not after I've sent it.
```

Why it works: the AI is uniquely good at this because it has no ego. It can switch from creator to critic instantly. Using this technique before any high-stakes deliverable is one of the highest-ROI habits you can build.

Quick Reference: Which Technique for Which Situation?

TECHNIQUE	BEST FOR	TRIGGER PHRASE
Chain-of-Thought	Complex reasoning, decisions	"Think step by step"
Few-Shot	Specific formats, style matching	"Here are 2 examples"
Persona Stacking	Multi-angle analysis, strategy	"Analyze from 3 perspectives"
Constraint Tightening	Reducing generic output	"Maximum X words, must/must not"
Prompt Chaining	Long documents, complex tasks	"Step 1: outline. Step 2: write"
Pre-Mortem	Quality control, risk assessment	"Attack this. What's weak?"

> **You don't need all six**
> Most professionals settle into 2-3 techniques that complement RCFC for their specific work. Chain-of-Thought and Constraint Tightening are the most universally useful. Few-Shot is essential for anyone who needs consistent formatting. Pre-Mortem is a game-changer for client-facing roles. Pick the ones that match your pain points, ignore the rest, and add more as you grow.

5 PROMPT PATTERNS YOU'LL USE EVERY DAY

The techniques above are powerful, but they're general-purpose — they tell you how to structure a prompt, not what to prompt for. What you need next are ready-made patterns: recurring prompt shapes that solve specific professional problems over and over again.

Think of it this way: RCFC and Chain-of-Thought are the grammar. The five patterns below are the sentences you'll actually say. Each one is a reusable template you can copy, adapt to your context, and deploy immediately. Together, they cover roughly 80% of what professionals need from an AI on any given workday.

> **Before you think "that's a lot of work for every prompt"**
> You might be reading this and wondering: do I really need to spell out Role, Context, Format, and Constraints every single time I open the chat? No, you don't. In the next chapter, you'll learn how to store your RCFC defaults inside dedicated projects and skills — so your role, your industry context, your preferred tone, and your recurring constraints are already loaded before you type a word. You write the full RCFC once; from then on, you only add what's specific to the task at hand.

Pattern 1: The Expert on Demand

Use when: you need deep expertise on a topic you're not specialized in.

```
You are [specific expert role with years of experience].
I am a [your role] working on [situation].
I need you to [specific task].
Explain your reasoning as you go.
Flag anything I should double-check with a real specialist.
```

Example:

```
You are a senior employment lawyer specializing in French
labor law with 20 years of experience.
I am a freelance consultant who just received a contract
```

```
from a new client that includes a non-compete clause.
Review the clause below and tell me:
- Is this enforceable under French law?
- What are the red flags?
- What modifications should I request?
Flag anything I should verify with an actual lawyer.

[paste clause here]
```

Why it works: the role anchors the AI's reasoning in domain expertise. The "flag anything" instruction prevents overconfidence and reminds you to verify critical points.

Pattern 2: The Format Transformer

Use when: you have content in one format and need it in another.

```
Transform the following [input format] into a [output format].
Keep the core message and key data points.
Adapt the tone for [target audience].
The output should be [length/structure constraints].

[paste content]
```

Example:

```
Transform these raw meeting notes into a structured
executive summary.
Keep all decisions and action items.
Adapt the tone for a C-level audience who wasn't present.
The summary should be under 300 words with sections:
Key Decisions, Action Items (with owners and deadlines),
and Open Questions.

[paste meeting notes]
```

Why it works: this is one of the highest-ROI uses of AI. You're not asking it to create from nothing — you're asking it to restructure what already exists. The quality of the input directly drives the quality of the output.

Pattern 3: The Sparring Partner

Use when: you need someone to challenge your thinking, find holes in your argument, or play devil's advocate.

```
I'm going to share [a plan / a proposal / an idea].
Your job is to challenge it constructively.
Find the 3 biggest weaknesses.
For each weakness, suggest a concrete fix.
Be direct — I don't need reassurance, I need improvement.

[paste your plan]
```

Example:

```
I'm about to propose a new pricing model to my client.
Your job is to challenge it constructively.
Find the 3 biggest weaknesses from the client's perspective.
For each weakness, suggest a concrete fix.
Be direct. I need this ready for a meeting tomorrow.

[paste pricing proposal]
```

Why it works: most professionals lack a sparring partner who can challenge their ideas on demand. The AI fills that role instantly. The key is the instruction to "be direct" — without it, the AI tends to be overly polite and bury its criticism in compliments.

Pattern 4: The Signal Extractor

Use when: you have a large amount of information and need to find the key insights.

```
Read the following [document / data / conversation].
Extract:
- The [N] most important [findings / risks / insights]
- Any [anomalies / contradictions / gaps]
- Recommended next steps based on the data
Rank by [impact / urgency / relevance to my situation].
Be concise. One paragraph per finding.

[paste content]
```

Example:

```
Read the following 50-page quarterly report.
Extract:
- The 5 most important financial findings
- Any data points that contradict management's narrative
- Recommended questions for the earnings call
Rank by impact on investment thesis.
Be concise. One paragraph per finding.

[paste or upload report]
```

Why it works: this turns a 2-hour reading task into a 10-minute review. The instruction to find contradictions is particularly powerful — the AI excels at cross-referencing different sections of a document to flag inconsistencies humans often miss under time pressure.

Pattern 5: The Constrained Writer

Use when: you need to produce written content that must follow specific rules, voice, or style.

```
Write a [type of content] about [topic].
Audience: [who will read this].
Tone: [specific tone descriptors].
```

```
Length: [word count or page count].
Structure: [sections, headings, format].
Must include: [specific points to cover].
Must avoid: [things NOT to include].
Reference style: [example of what “good” looks like, or
paste a sample of your own writing].
```

Example:

```
Write a LinkedIn post about the importance of data quality
in ERP migrations.
Audience: IT directors and CFOs in manufacturing.
Tone: authoritative but approachable, no jargon.
Length: 150-200 words.
Structure: hook + 3 key points + call to action.
Must include: a real-world consequence of bad data quality.
Must avoid: buzzwords like “digital transformation” or
“leveraging AI.” No hashtags.
Reference style: write like the posts I’ve been sharing
on my profile (conversational, direct, opinionated).
```

Why it works: the “must avoid” section is where most of the value lives. Without it, the AI defaults to the most common patterns in its training data — which are exactly the clichés you want to avoid. The reference style instruction is especially powerful if you paste an actual sample of your writing.

Quick Reference: Which Pattern for Which Situation?

PATTERN	BEST FOR	EXAMPLE TASKS
Expert on Demand	Getting domain-specific analysis	Legal review, tax advice, technical audit
Format Transformer	Restructuring existing content	Notes → summary, email → slides, raw data → report
Sparring Partner	Stress-testing your ideas	Proposal review, strategy critique, risk assessment
Signal Extractor	Finding insights in large data	Report analysis, contract review, market research

Constrained Writer	Producing content under rules	LinkedIn posts, client emails, proposals, reports

YOUR FIRST SKILL: THE RCFC PROMPT BUILDER

This is where you move from "reading about prompting" to "having a tool that does it for you." The RCFC Prompt Builder is a skill you install in Claude (or adapt as a Custom GPT in ChatGPT / Gem in Gemini) that automatically transforms any rough prompt into a structured, professional-grade prompt.

How It Works

1. **You type your rough request** in plain language: "I need to write an email to decline a meeting."
2. **The skill asks you 4-6 targeted questions** to fill in the RCFC elements: Who are you in this context? Who's the recipient? What tone? What should you avoid saying? What's the format?
3. **The skill generates a fully structured prompt** using the RCFC method, incorporating all your answers.
4. **It then immediately executes the prompt** and delivers the output.

Over time, you'll internalize the RCFC method and stop needing the skill for simple tasks. But for complex or high-stakes prompts, you'll keep coming back to it.

The Skill File: RCFC Prompt Builder

Copy the following text into a file and add it to your Claude Project (or adapt it as Custom GPT instructions for ChatGPT). The skill will activate whenever you type a trigger phrase like "RCFC", "optimize my prompt", or "help me write a better prompt."

```
---
name: rcfc-prompt-builder
description: >
  Transforms a rough user prompt into a structured,
  professional-grade prompt using the RCFC method
```

```
  (Role, Context, Format, Constraints). Works in two
  phases: (1) asks targeted questions to fill RCFC
  gaps, (2) generates and executes the optimized prompt.
---

# RCFC Prompt Builder

## Purpose
Transform any rough, incomplete, or vague prompt into
a structured, high-quality prompt using the RCFC
framework (Role, Context, Format, Constraints), then
execute it immediately.

## Trigger phrases
Activate this skill when the user says any of:
- "RCFC"
- "optimize my prompt"
- "help me write a better prompt"
- "improve this prompt"
- "structure my request"
- "use the RCFC method"

## Workflow

### Phase 1: Analyze and Ask

1. Read the user's rough prompt or task description.
2. Identify which RCFC elements are missing or weak:
   - R (Role): Is there a clear expert perspective?
   - C (Context): Is the situation fully described?
     Who is the audience? What's the background?
   - F (Format): Is the desired output specified?
     Length, structure, tone, style?
   - C (Constraints): Are there boundaries?
     Things to include, exclude, or avoid?
3. Ask ONLY the missing questions. If the user already
   provided a clear role, don't ask about role.
   Maximum 6 questions, minimum 2.
4. Present the questions in a numbered list.
```

```
5. STOP and wait for answers.

### Phase 2: Build and Execute

Once the user answers:
1. Construct the optimized prompt using this structure:

   **Role:** [Expert role with relevant experience]
   **Context:** [Full situation, audience, background]
   **Format:** [Output type, length, structure, tone]
   **Constraints:** [Boundaries, inclusions, exclusions]

2. Show the optimized prompt to the user in a
   clearly labeled block.
3. Ask: "Should I execute this prompt now, or would
   you like to adjust anything first?"
4. If the user says execute (or yes / go / do it),
   run the prompt and deliver the output.
5. If the user wants adjustments, modify the prompt
   and ask again.

## Question Bank

Draw from these when identifying missing RCFC elements:

### Role questions:
- What kind of expert should handle this task?
- What level of seniority or specialization is needed?

### Context questions:
- Who is the target audience for this output?
- What's the situation or background I should know?
- What have you already tried or decided?
- Is there an existing document or example I should
  reference?

### Format questions:
- What format do you want? (email, report, bullets,
  presentation outline, table, etc.)
```

```
- How long should the output be?
- What tone are you going for? (formal, conversational,
  direct, empathetic, etc.)

### Constraint questions:
- Anything you want me to specifically include?
- Anything I should definitely avoid?
- Any company-specific terminology or style rules?

## Rules
- Never ask more than 6 questions.
- Never ask about elements the user already provided.
- Always show the constructed prompt before executing.
- Use the same language as the user.
- If the user's request is already well-structured
  (all 4 RCFC elements present), say so and offer
  to execute directly or suggest minor improvements.
```

Installation Guide

For Claude:

- Copy the skill text above.
- Open a new Claude conversation and paste it with a message like: "Turn this into a proper skill file I can download."
- Claude will generate a clean skill file. Download it.
- In Claude, click "Customize" (bottom left) → "Skills" → "Create skill" → "Upload skill."
- Upload the file. The skill is now available in all your conversations — type "RCFC" or "optimize my prompt" to trigger it.

For ChatGPT:

- Go to "Explore GPTs" → "Create."
- Paste the skill content into the Instructions field.
- Name it "RCFC Prompt Builder."

- Use it by opening the GPT and typing your rough prompt.

For Gemini:

- Go to Gemini → Gems → "Create a Gem."
- Paste the skill content into the instructions.
- Save and activate.

Example: The RCFC Skill in Action

You type:

```
RCFC: I need to write a LinkedIn post about why
companies fail at ERP migrations.
```

The skill responds:

```
I can see your topic but I'm missing a few things.
Let me ask:

1. What's your professional role? (This helps me set
   the right expert perspective.)
2. Who's your target audience on LinkedIn?
   (IT directors? CFOs? Consultants?)
3. What tone do you prefer? (Provocative? Educational?
   Conversational?)
4. Any specific angle or example you want to include?
5. Anything to avoid? (Buzzwords, hashtags,
   company mentions?)
```

You answer:

```
1. SAP FICO consultant, 10 years experience
2. IT directors and project managers in manufacturing
3. Direct, opinionated, a bit provocative
4. Data migration disasters I've personally seen
5. No hashtags, no "digital transformation" buzzwords
```

The skill builds and shows the optimized prompt, then executes it.

> **The learning effect**
> After using the RCFC skill 10-15 times, you'll notice something: you start writing RCFC prompts naturally, without the skill. The questions become automatic. That's the point. The skill is a training tool as much as a productivity tool. Eventually, you graduate from it — and that's a success, not a loss.

PART 3

Build Your Persistent Context System

Including the AI Onboarding Skill

The Professional's AI Playbook

STOP RE-INTRODUCING YOURSELF TO YOUR AI

Here's the single biggest unlock most professionals miss: you can onboard your AI the same way you'd onboard a new colleague. You do it once, and from that point on, every conversation starts with shared context instead of a blank slate.

Think about what happens when a new team member joins your company. On day one, they know nothing about your clients, your processes, your writing style, or your internal vocabulary. So you brief them. You share documents. You explain how things work. After a week, they're productive. After a month, they anticipate your needs.

An AI chatbot works exactly the same way — except nobody told you to onboard it. So you've been working with a colleague who has amnesia every single morning.

This chapter fixes that. You're going to create five context files that tell the AI who you are, how you work, and what good output looks like. Once set up, these files load automatically in every conversation within your project. No more cold starts. No more "I'm a consultant who works in..." at the top of every chat.

> **This is where RCFC becomes effortless**
> Remember the RCFC method from Part 2? Your context files permanently store the R (your default role), the first C (your ongoing context), and many of your recurring Constraints. From now on, you only need to specify the Format and task-specific details. The heavy lifting is done once.

THE 5 CONTEXT FILES: YOUR AI'S ONBOARDING KIT

Each file serves a distinct purpose. Together, they form a comprehensive briefing that makes the AI understand your professional world from the first message.

File 1: Professional Identity Card

Purpose: Tell the AI who you are before every task.

What to include:

- Your job title and area of expertise
- Your company or organization (size, industry, structure)
- Your current priorities and active projects
- Your clients or stakeholders (types, not names if sensitive)
- Your level of seniority and decision-making authority

Template:

```
# Professional Identity Card

## Role
I am a [job title] with [X] years of experience in
[industry/domain]. I specialize in [specific areas].

## Organization
I work at/for [company type: startup, consulting firm,
multinational, freelance]. The company operates in
[industry] with [size: headcount or revenue bracket].
[If freelance: I operate through [legal structure] and
bill [daily/hourly rate bracket if relevant]].

## Current Priorities
- [Priority 1: e.g., managing an ERP migration project]
- [Priority 2: e.g., developing a new service offering]
```

```
- [Priority 3: e.g., building an AI-powered workflow]

## Clients/Stakeholders
My typical clients are [description: CFOs of mid-size
manufacturers, IT directors in pharma, etc.].
I report to / collaborate with [key stakeholders].

## Seniority
[Senior / mid-level / junior]. I [do/don't] have
direct reports. I [am/am not] the final decision-maker
on [type of decisions].
```

File 2: Communication Style Guide

Purpose: Ensure every output matches the way you actually write and communicate.

What to include:

- Your preferred tone (formal, conversational, direct, diplomatic)
- Your writing habits (short sentences? long paragraphs? bullet points?)
- Words or phrases you love or hate
- Formatting preferences (headers, bold, numbered lists)
- A sample of your actual writing (paste 2-3 paragraphs you're proud of)

Template:

```
# Communication Style Guide

## Tone
Default tone: [direct and analytical / warm and
empathetic / authoritative but approachable / etc.].
Adjust based on audience:
- Clients: [more formal / more diplomatic]
- Internal team: [more casual / more direct]
- Public (LinkedIn, articles): [more opinionated /
```

```
  more educational]

## Writing Style
- I prefer [short sentences / concise paragraphs].
- I [use / avoid] bullet points.
- I [use / avoid] headers and sub-headers.
- My paragraphs are typically [2-3 / 4-5] sentences.

## Language Preferences
Words I like: [e.g., “concretely,” “actionable,”
“the key takeaway is”]
Words I avoid: [e.g., “leverage,” “synergy,”
“digital transformation,” “in today’s fast-paced
world”]

## Writing Sample
[Paste 2-3 paragraphs of your best writing here.
This gives the AI a concrete reference for your voice.]
```

File 3: Industry Glossary

Purpose: Ensure the AI uses your professional vocabulary correctly, including internal jargon, acronyms, and domain-specific terms.

What to include:

- Industry-specific acronyms and their definitions
- Internal company terms (project names, process names, tool names)
- Terms that have a specific meaning in your context vs. their common meaning
- Competitor or partner names the AI should know

Template:

```
# Industry Glossary

## Acronyms
- [ABC]: [what it stands for and what it means]
- [DEF]: [definition]
- [GHI]: [definition]

## Internal Terms
- [Project Phoenix]: [what it is]
- [The Bridge]: [internal name for X process]
- [Tiger Team]: [what it means in our context]

## Domain-Specific Definitions
- [Term]: In our context, this means [specific
  definition], not the common meaning of [X].

## Key Players
- [CompanyA]: our main competitor in [segment]
- [CompanyB]: our technology partner for [X]
- [ToolName]: the platform we use for [Y]
```

File 4: Deliverable Templates Library

Purpose: Give the AI examples of your actual deliverables so it can replicate your format and standards.

What to include:

- The 3-5 types of documents you produce most often
- For each: the standard structure (sections, headings, length)
- Optionally: a redacted example of each

Template:

```
# Deliverable Templates

## Template 1: [e.g., Client Status Report]
Structure:
1. Executive Summary (3-5 sentences)
2. Progress This Period (bullet points by workstream)
3. Risks and Issues (table: risk, impact, mitigation)
4. Next Steps (numbered, with owners and dates)
5. Appendix (if needed)
Length: 1-2 pages
Tone: factual, concise, no filler

## Template 2: [e.g., Client Proposal]
Structure:
1. Context and Understanding
2. Proposed Approach
3. Deliverables and Timeline
4. Team and Expertise
5. Investment
Length: 5-8 pages
Tone: confident, solution-oriented

## Template 3: [e.g., Meeting Minutes]
[same format as above]
```

File 5: Preferences Journal

Purpose: A living document where you record what works and what doesn't with the AI, so it continuously improves.

What to include:

- Prompt patterns that consistently produce great results
- Things the AI tends to get wrong that you need to correct
- Output preferences you've discovered through experience
- Recurring instructions you're tired of typing

Template:

```
# Preferences Journal

## What Works Well
- When I ask for [X], always [approach/format]
- For emails, always include a clear subject line
- When analyzing documents, start with a 3-sentence
  summary before diving into details

## Common Mistakes to Avoid
- Don't use the word "delve" — ever
- Don't add motivational conclusions ("In conclusion,
  by following these steps you'll achieve...")
- Don't use exclamation marks in professional content
- Don't add caveats like "I'm an AI and can't..."

## Standing Instructions
- Always default to [metric system / imperial]
- Currency: [EUR / USD / GBP]
- Date format: [DD/MM/YYYY]
- When I say "short," I mean under [150] words
- When I say "detailed," I mean [500-800] words

## Last Updated: [date]
```

The 5 Files at a Glance

FILE	WHAT IT COVERS	HOW OFTEN TO UPDATE
Identity Card	Who you are, what you do, who you serve	When your role or priorities change
Style Guide	How you write, your tone, your voice	Rarely — your style is stable
Glossary	Your professional vocabulary and jargon	When you enter a new domain or project
Templates	Your standard deliverable formats	When you create a new type of deliverable
Preferences	What works, what doesn't, standing rules	Continuously — after every "aha" moment

TWO WAYS TO CREATE YOUR CONTEXT FILES

You have two options. Both produce the same result. Pick whichever fits your style.

Path A: The Guided Path — Use the AI Onboarding Skill

Launch the skill, answer its questions, and it generates all five files for you. This is the fastest route if you want to get set up in 20 minutes without staring at blank templates. (Don't worry if you haven't installed a skill before — full installation steps for Claude, ChatGPT, and Gemini are provided at the end of this chapter.)

The skill adapts its questions to your profile. A freelance consultant gets different questions than a corporate manager. A finance professional gets different prompts than a marketing lead.

Path B: The Manual Path — Fill the Templates Yourself

Copy the templates from the previous section into your preferred tool (Google Docs, Notion, or directly into your AI platform). Fill them in at your own pace. This works well if you want full control over every word or if you prefer to think things through slowly.

> **Our recommendation**
> Start with Path A (the skill) to generate a first draft in minutes. Then review and refine manually. You get the speed of automation plus the precision of your own judgment. Best of both worlds.

WHERE TO PUT YOUR CONTEXT FILES

Once your five files are ready, you need to load them into your AI platform so they're available automatically. Here's how to do it on each platform.

Claude: Projects

1. Go to Claude.ai and click "Projects" in the sidebar.
2. Create a new project (e.g., "My Professional Workspace").
3. Click "Add content" and upload or paste each of your 5 files.
4. Optionally, add project-level instructions in the "Custom instructions" field (e.g., "Always read and apply the context files before responding").
5. Every conversation you start inside this project will automatically have access to all five files.

Pro tip: create multiple projects for different areas of your work (e.g., "Client A," "Business Development," "Internal Admin"). Share the Identity Card and Style Guide across all of them, but customize the Glossary and Templates per project.

ChatGPT: Custom GPTs

1. Go to "Explore GPTs" → "Create."
2. In the "Instructions" field, paste the content of your 5 files (or the most critical ones — Identity Card + Style Guide + Preferences at minimum).
3. Use the "Knowledge" section to upload larger files (glossary, templates, reference documents).
4. Name and save your GPT. Use it as your default starting point for work conversations.

Alternative: if you don't want to create a custom GPT, you can paste the essential context (Identity Card + Preferences) into Settings → Custom Instructions. This applies globally to all conversations but is limited in length.

Gemini: Gems

1. Open Gemini and go to the Gems section.
2. Create a new Gem.
3. Paste your context files into the Gem instructions.
4. Save and use this Gem as your professional workspace.

Gemini advantage: if you're in the Google ecosystem, Gemini can also pull context from your Gmail, Drive, and Calendar via its built-in integrations. Your context files plus Google integrations create a particularly powerful combination.

YOUR SECOND SKILL: THE AI ONBOARDING ASSISTANT

This skill turns the potentially tedious process of writing five context files into a guided, conversational experience. Launch it, answer its questions, and it generates all five files for you — ready to copy into your AI platform.

How It Works

1. **You trigger the skill** by saying "onboard me," "set up my context," or "create my profile."
2. **The skill asks you a series of structured questions** — about your role, your industry, your writing style, your clients, your recurring deliverables, and your preferences.
3. **Based on your answers, it generates all 5 context files** pre-filled with your information.
4. **You review, edit if needed, and install them** in your project or custom GPT.

The skill adapts to your profile. It detects whether you're a consultant, a manager, a freelancer, or a specialist, and adjusts its questions accordingly.

The Skill File: AI Onboarding Assistant

Copy the following into a file and add it to your Claude Project (or adapt for ChatGPT / Gemini).

```
---
name: ai-onboarding-assistant
description: >
  Interactive skill that interviews the user about
  their professional context and generates 5 ready-to-use
  context files: Identity Card, Style Guide, Industry
  Glossary, Deliverable Templates, and Preferences
  Journal. Adapts questions to the user's detected
  profile (consultant, manager, freelancer, specialist).
---

# AI Onboarding Assistant
```

```
## Purpose
Interview the user about their professional life and
generate 5 structured context files they can install
in their AI workspace for persistent personalization.

## Trigger phrases
Activate when the user says any of:
- "onboard me"
- "set up my context"
- "create my profile"
- "generate my context files"
- "AI onboarding"
- "set up my AI workspace"
- "help me configure my AI"

## Workflow

### Phase 1: Quick Profile Detection

Ask one opening question:
"In one or two sentences, describe what you do
professionally — your role, your industry, and whether
you work in a company, as a freelancer, or something
else."

Based on the answer, classify the user as one of:
- CONSULTANT (external advisor, client-facing)
- MANAGER (team lead, department head, executive)
- FREELANCER (independent, solo operator)
- SPECIALIST (individual contributor, deep expertise)

This classification determines which questions to
prioritize in Phase 2.

### Phase 2: Structured Interview

Ask questions in 3 batches (to avoid overwhelming
the user). Wait for answers between each batch.
```

```
**Batch 1 — Identity & Role (4-5 questions):**
- What is your exact job title?
- How many years of experience do you have?
- What are your 2-3 areas of deepest expertise?
- Who are your typical clients or stakeholders?
  (If CONSULTANT/FREELANCER: client types.
  If MANAGER/SPECIALIST: internal stakeholders.)
- What are your top 3 professional priorities right
  now?

**Batch 2 — Communication & Style (4-5 questions):**
- How would you describe your communication style
  in 3 words?
- Do you prefer bullet points or flowing prose?
- Are there any words or phrases you love using?
- Are there any words or clichés you absolutely
  hate?
- Can you paste 2-3 sentences from something you
  wrote recently that represents your style well?
  (Optional but very valuable.)

**Batch 3 — Work Patterns & Preferences
  (4-5 questions):**
- What are the 3-5 types of documents or deliverables
  you produce most often?
- For each, briefly describe the standard structure
  or format.
- What are the most important industry-specific
  terms or acronyms someone working with you needs
  to know? (List 5-10.)
- What tools do you use daily?
  (Email client, project management, CRM, etc.)
- Think about the last time an AI gave you a bad
  output. What was wrong with it? (This helps build
  the Preferences Journal.)

### Phase 3: Generate the 5 Files

Once all batches are answered, generate the 5 files
```

```
using the following structure:

**File 1: Professional Identity Card**
Sections: Role, Organization, Current Priorities,
Clients/Stakeholders, Seniority.
Populate from Batch 1 answers.

**File 2: Communication Style Guide**
Sections: Tone, Writing Style, Language Preferences,
Writing Sample.
Populate from Batch 2 answers.

**File 3: Industry Glossary**
Sections: Acronyms, Internal Terms,
Domain-Specific Definitions, Key Players.
Populate from Batch 3 + infer common acronyms
from the user's stated industry.

**File 4: Deliverable Templates Library**
For each deliverable type the user listed:
structure, length, tone.
Populate from Batch 3 answers.

**File 5: Preferences Journal**
Sections: What Works Well, Common Mistakes to Avoid,
Standing Instructions.
Populate from Batch 2 + Batch 3 answers.
Pre-fill standing instructions with date format,
currency, and measurement system based on detected
locale.

### Phase 4: Deliver and Guide

1. Present all 5 files clearly, each in its own
   labeled section.
2. After each file, add a note: "Review this file.
   Want me to adjust anything before we move to
   the next one?"
3. Once all files are approved, provide platform-
```

```
   specific installation instructions:
   - Claude: add to Project content
   - ChatGPT: paste in Custom GPT instructions +
     Knowledge
   - Gemini: paste in Gem instructions
4. Offer: "Would you like me to also generate a
   single consolidated file you can use in
   system-level instructions for quick setup?"

## Rules
- Never ask more than 5 questions per batch.
- Always wait for answers between batches.
- If the user seems impatient, offer to skip to
  file generation with what you have.
- Use the same language as the user.
- When inferring glossary terms, be conservative —
  include only terms you're confident about and
  ask the user to validate.
- The generated files should be ready to copy-paste
  as-is, but invite the user to personalize them
  further.
```

Installation Guide

For Claude:

- Copy the skill text above and paste it into a Claude conversation with: "Turn this into a proper skill file I can download."
- Download the generated file.
- Go to Customize → Skills → Create skill → Upload skill.
- Start a conversation and type "onboard me."
- Answer the 3 batches of questions.
- Review the generated context files, then add them to your Projects as content.

For ChatGPT:

- Create a Custom GPT with the skill instructions.
- Run the onboarding, then copy the generated files into a separate Custom GPT for daily use (or paste them into Custom Instructions).

For Gemini:

- Create a Gem with the skill instructions.
- Run the onboarding, then transfer the generated files into your primary workspace Gem.

Example: The Onboarding Skill in Action

You type:

```
Onboard me.
```

The skill responds:

```
Welcome! I'm going to help you create your AI
context files. This takes about 15-20 minutes and
will save you hours every week.

Let's start with the big picture. In one or two
sentences, describe what you do professionally —
your role, your industry, and whether you work in
a company, as a freelancer, or something else.
```

You answer:

```
I'm a freelance management consultant with 12 years
of experience. I specialize in operations and process
improvement for mid-size manufacturing companies.
I operate as a solo consultant through my own company.
```

The skill detects: FREELANCER/CONSULTANT profile, then asks Batch 1:

```
Great — I've got a clear picture. You're a freelance
```

```
consultant in manufacturing operations. Let me ask
a few more things to build your Identity Card.

1. What are your 2-3 deepest areas of expertise
   within operations? (e.g., lean manufacturing,
   supply chain, ERP implementation)
2. Describe your typical client in one sentence.
   (e.g., "COO of a 200-500 person manufacturer")
3. What are your top 3 priorities right now?
4. What's the legal structure of your business?
   (Relevant for how you describe your setup.)
5. What's your seniority level when embedded in
   a client's team? (Advisory? Hands-on? Both?)
```

The conversation continues through Batches 2 and 3, then generates all 5 files.

> **The compound effect**
> These context files don't just improve individual conversations — they compound over time. As you update your Preferences Journal and refine your templates, the AI gets better at anticipating your needs. After 30 days of use, you'll wonder how you ever worked without them.

PART 4

Produce Visual Deliverables Directly from the Chat

Presentations, Documents, Diagrams & More

The Professional's AI Playbook

YOUR AI CAN MAKE MORE THAN TEXT

Most professionals think of AI chatbots as text-only tools: you type a question, you get a written answer. But in 2026, every major platform can produce polished visual deliverables — presentations, formatted documents, diagrams, and charts — directly from the conversation.

This is a genuine game-changer for anyone who spends hours in PowerPoint, Word, or Google Slides every week. Instead of manually building slides, formatting reports, or drawing diagrams, you describe what you need and the AI creates it.

The catch: each platform handles this differently. Claude, ChatGPT, and Gemini each have distinct strengths and workflows for visual deliverables. This chapter maps them out so you know exactly which tool to reach for and how to use it.

CREATING PRESENTATIONS WITH AI

This is the deliverable professionals ask about most. Here's how each platform handles it, with honest assessments of what works and what doesn't.

Claude: Built-in .pptx Generation

Claude can generate complete PowerPoint files (.pptx) directly in the chat. You describe what you need, and Claude creates a downloadable file with real slides, text, layouts, and formatting. No plugins, no third-party tools, no export steps.

How to do it:

1. Open a conversation in Claude (ideally inside a Project with your context files loaded).
2. Describe your presentation: topic, audience, number of slides, structure, and tone.
3. Claude generates the .pptx file and provides a download link.
4. Download, review, and refine. Ask Claude to adjust specific slides if needed.

Example prompt:

```
Create a 10-slide PowerPoint presentation for a client
quarterly business review.

Audience: CFO and operations director of a mid-size
manufacturing company.
Tone: professional, data-driven, confident.

Structure:
1. Title slide with our company name and date
2. Agenda
3. Executive summary (3 key takeaways)
4-6. Progress by workstream (3 slides)
7. Key metrics dashboard (table format)
8. Risks and mitigation
9. Next quarter priorities
```

```
10. Q&A / contact info

Use a clean, modern design. Dark blue and white color
scheme. No clip art.
```

Strengths: fully self-contained, no third-party dependency, produces a real .pptx file you can open in PowerPoint or Google Slides, great at following structural instructions.

Limitations: design quality is functional rather than stunning — think "clean corporate template" not "award-winning design." Best used to generate the structure and content, then refine the design manually or with your company template.

ChatGPT: Three Ways to Create Presentations

ChatGPT offers three workflows for presentations, from quickest to most polished:

Option 1: Agent Mode — Direct .pptx Generation

ChatGPT can generate .pptx files directly in the chat using Agent mode. This is the fastest path: you describe your presentation, and ChatGPT creates a downloadable PowerPoint file. Available to Plus, Pro, and Team subscribers.

1. Open a new chat in ChatGPT.
2. Select "Agent" or go to Tools → Agent mode in the lower left corner.
3. Describe your presentation (topic, audience, number of slides, structure).
4. ChatGPT generates the .pptx file and provides a download link.

Strengths: fast, no third-party tool needed, produces a real .pptx you can open in PowerPoint or Google Slides.

Limitations: the design is basic — clean but minimal. It can't apply your corporate template, and layout options are limited. Best used for a quick first draft you'll polish manually.

Option 2: ChatGPT + Canva (best design quality)

ChatGPT integrates directly with Canva. You can create, preview, and edit Canva presentations without leaving the chat. The results are visually polished, with access to Canva's full template library.

- In ChatGPT, go to Settings → Apps → Add Canva, and connect your Canva account.
- In your chat, mention Canva in your prompt: "Canva, create a 10-slide presentation about..."
- ChatGPT generates the slides and shows a preview. Edit through conversation or click "Open in Canva" to fine-tune.
- Export from Canva as .pptx, PDF, or share directly.

Strengths: beautiful designs, access to Canva's templates and brand kits, can edit text across all slides at once through conversation.

Limitations: requires a Canva account (free tier works for basics, Pro for brand kits), adds a dependency on a third-party platform.

Option 3: ChatGPT as Content Engine + Manual Design

Ask ChatGPT to generate the full content and structure for each slide (title, bullet points, speaker notes), then build the slides yourself in PowerPoint or Google Slides. This works well when you have a company template you must use.

```
Create a detailed outline for a 12-slide investor pitch
deck. For each slide, provide:
- Slide title
- 3-4 bullet points of content
- Speaker notes (what I should say, 50-75 words)
- Suggested visual element (chart type, image idea, etc.)

Structure: Problem, Solution, Market size, Product demo,
Business model, Traction, Team, Financials, Ask.
```

Gemini: Native Google Slides Integration

Gemini has the tightest integration with Google Slides, which makes it the natural choice if your company runs on Google Workspace.

Option 1: Generate from the Gemini App (Canvas)

1. Open the Gemini web app.
2. Click Tools → Canvas at the bottom of the chat box.
3. Type your prompt: "Build a presentation about..." with detailed slide descriptions.
4. Gemini generates a full presentation preview in the chat.
5. Click "Export to Slides" to save it as a Google Slides file in your Drive.

Option 2: Gemini Sidebar Inside Google Slides

If you're already working in Google Slides, the Gemini sidebar can generate individual slides, write content, create images, and reference your Drive files — all without leaving the app.

1. Open a presentation in Google Slides.
2. Click "Ask Gemini" (top right) to open the sidebar.
3. Type prompts like "Create a slide summarizing the Q3 report from @quarterly-report.docx"
4. Gemini generates the slide; click "Insert" to add it to your deck.

Strengths: seamless Google Workspace integration, can pull data from Drive files and Gmail directly into slides, fully editable output that matches your existing theme, generates images and diagrams inline.

Limitations: currently generates one slide at a time in the sidebar (Canvas generates full decks). Requires a Google AI Pro or Workspace Business+ subscription for full features.

Presentation Creation: Platform Comparison

	CLAUDE	CHATGPT	GEMINI
Output format	.pptx file	.pptx (Agent) / Canva	Google Slides
Full deck at once	Yes	Yes (Agent or Canva)	Yes (via Canvas)

Design quality	Functional	Basic (Agent) / High (Canva)	High (Google themes)
Third-party needed	No	No (Agent) / Canva (design)	No (Google ecosystem)
Iteration in chat	Strong	Strong	Moderate
File references	Upload to Project	Upload or Drive link	Drive + Gmail native
Best for	Self-contained .pptx	Beautiful design	Google Workspace users

BEYOND SLIDES: OTHER VISUAL DELIVERABLES

Presentations get the most attention, but AI can also produce several other types of visual outputs that professionals use daily.

Formatted Documents (Word, PDF)

Claude can generate .docx (Word) and .pdf files directly from the chat. This is ideal for formatted reports, proposals, letters, and any document that needs to look professional when shared.

```
Create a Word document for a project status report.
Include a header with our project name and date.
Use the structure from my Deliverable Templates (File 4).
Populate it with the following updates:
[paste your raw status notes]
```

ChatGPT can produce formatted content via Canvas (its built-in document editor) or export through plugins. Gemini writes content that you can export directly to Google Docs.

Diagrams and Flowcharts

All three platforms can generate diagrams using Mermaid syntax (a text-based diagramming language). Claude also renders diagrams as visual artifacts directly in the chat.

```
Create a flowchart showing our client onboarding process:
1. Lead qualification
2. Discovery call
3. Proposal sent
4. Negotiation (branch: accepted / rejected)
5. If accepted: contract signing → kickoff meeting
6. If rejected: feedback collection → archive

Use a clean, professional style.
```

This is particularly powerful for process documentation, architecture diagrams, org charts, and decision trees — deliverables that usually require Visio, Lucidchart, or hours in PowerPoint.

Charts and Data Visualizations

All three platforms can create charts from data you provide:

- **Claude:** generates interactive charts as artifacts (HTML/JavaScript) directly in the chat, or embeds charts in .pptx and .docx files.
- **ChatGPT:** uses Code Interpreter to generate charts with Python (matplotlib, plotly), exports as images or interactive HTML.
- **Gemini:** creates charts in Google Sheets or directly in Slides via the sidebar.

```
Here's our quarterly revenue data:
Q1: 450K, Q2: 520K, Q3: 480K, Q4: 610K

Create a clean bar chart with a trend line.
Use dark blue for the bars.
Add percentage change labels between quarters.
Title: "Revenue Growth FY2025"
```

Images and Visual Assets

Need a quick visual for a presentation, a social post, or a document illustration? ChatGPT and Gemini can generate images natively:

- **ChatGPT:** uses GPT Image (native generation) — excellent text rendering, photorealistic output, iterative editing through conversation.
- **Gemini:** uses Nano Banana (native generation) — fast, multiple style options, integrated with Google ecosystem. Can even use your Google Photos as references.
- **Claude:** does not generate images. Use it for the content and structure, then generate visuals with ChatGPT or Gemini.

THE UNIVERSAL CREATIVE BRIEF PROMPT

Here's a reusable prompt template that works for any visual deliverable — presentations, documents, diagrams, or charts. It uses the RCFC method you learned in Part 2, adapted specifically for visual outputs.

```
CREATIVE BRIEF

Deliverable type: [presentation / document / diagram /
  chart / infographic]

Audience: [who will see this and what do they care about]

Purpose: [what should the viewer think, feel, or do
  after seeing this]

Content outline:
  - [Section 1: title + key points]
  - [Section 2: title + key points]
  - [Section 3: title + key points]
  (add as many sections as needed)

Design direction:
  - Color scheme: [specific colors or general vibe]
  - Style: [minimal / corporate / creative / playful]
  - Visual elements: [charts, icons, photos, diagrams]

Constraints:
  - Length: [number of slides / pages / dimensions]
  - Brand: [company name, logo notes, font preferences]
  - Must include: [specific data points, legal text,
    disclaimer]
  - Must avoid: [stock photos, clip art, specific words]

Raw content: [paste your notes, data, or reference
  material below]
```

The 80/20 of visual deliverables
Let the AI handle the 80%: structure, content, initial layout, and formatting. You handle the 20%: brand alignment, final design tweaks, and quality review. This workflow typically saves 60-70% of the time you'd spend building deliverables from scratch, while producing results that are as good or better than what you'd create manually under time pressure.

Quick Reference: Which Tool for Which Deliverable?

DELIVERABLE	CLAUDE	CHATGPT	GEMINI
PowerPoint / Slides	**Best (.pptx native)**	Good (via Canva)	**Best (Google Slides)**
Word documents	**Best (.docx native)**	Good (Canvas)	Good (export to Docs)
PDF reports	**Best (.pdf native)**	Moderate	Moderate
Diagrams / flowcharts	**Best (artifacts)**	Good (code)	Good (Slides sidebar)
Charts / data viz	Strong (artifacts)	Strong (Code Interpreter)	Strong (Sheets)
Images / visual assets	No	**Best (GPT Image)**	Strong (Nano Banana)
Spreadsheets	Good (.xlsx native)	Strong (Code Interpreter)	**Best (Google Sheets)**

The power combo
Many professionals use two tools together: Claude for structure, content, and file generation (documents, presentations, analysis) plus ChatGPT or Gemini for images and visual polish. This combination covers every deliverable type without compromise.

From the author: how I compensate for my lack of design skills
I'm not a designer. I have no eye for color palettes, no instinct for layout, and my PowerPoint slides used to look like they were made in 2007. But here's how I compensate: I use the few-shot prompting technique from Part 2.

Instead of trying to explain a theme, a color scheme, or a visual style in words — which I'm terrible at — I upload an example. A PowerPoint I liked from a colleague. A Word document with formatting I want to replicate. A slide deck from a client that set the bar. Then I simply tell

> the AI: “Use this as a template for the next document” or “Describe this design so I can reuse the style later.”
>
> This is a perfect example of where AI truly has your back. You don’t need to become a designer. You just need one good example and the right prompt. The AI bridges the gap between “I know what looks good when I see it” and “I can produce it myself.”

PART 5

10 Ready-to-Use Professional Playbooks

Copy, Paste, Adapt, Ship

The Professional's AI Playbook

YOUR PROFESSIONAL PLAYBOOK LIBRARY

Each playbook on the following pages is a complete, self-contained workflow you can use immediately. They're built with the RCFC method from Part 2, designed to work with the context files from Part 3, and ready to copy-paste into any AI chatbot.

Every playbook follows the same format:

- When to use it — the situation that triggers this playbook.
- The prompt — a complete, ready-to-use prompt. Copy it verbatim and fill in the brackets.
- What you get — the expected output so you know what good looks like.
- Variants — alternative versions for different contexts or audiences.

> **Warning: these prompts will fall flat without context**
> If you copy-paste one of these playbooks into a blank conversation with no context, the output will be generic, vague, and ultimately useless. The AI doesn't know your industry, your clients, your tone, or your standards — so it fills every gap with defaults. That's the "this tool doesn't work for me" moment we want you to avoid.
>
> **You have two options to make these prompts work:**
>
> Option A: Use them inside a Project (Claude), Custom GPT (ChatGPT), or Gem (Gemini) where your context files from Part 3 are already loaded. The AI knows you before you type a word, and the playbook just adds the task-specific instructions on top.
>
> Option B: If you're in a plain conversation, paste the relevant context directly into the prompt — your role, your audience, the background, a reference document. Fill every bracket in the prompt with real information. The more specific you are, the more usable the output.

Playbook 1: Prepare a Client Meeting in 10 Minutes

When to use it:

You have a meeting in an hour and you haven't had time to prepare properly. You need a briefing sheet with the key talking points, potential client questions, and your recommended positions.

The prompt:

```
Role: You are my strategic advisor preparing me for a
client meeting.

Context:
- Client: [client name/description]
- Meeting topic: [what we're discussing]
- Background: [paste any relevant context — last
  meeting notes, email thread, project status]
- My goal for this meeting: [what I want to achieve]

Create a meeting prep sheet with:
1. 3-5 key talking points (with supporting arguments)
2. 5 questions the client is likely to ask, with
   suggested answers
3. 2-3 potential objections and how to handle them
4. Recommended meeting structure (agenda, 30 or 60 min)
5. One killer opening statement to set the tone

Keep it concise. I'll read this in 10 minutes.
```

What you get:

A structured, one-page briefing sheet you can review on your phone on the way to the meeting. It covers your talking points, anticipated questions, and a recommended flow.

Variants:

- **High-stakes negotiation:** Add "include BATNA analysis and walk-away conditions"
- **Recurring status meeting:** Add "focus on what changed since last meeting and decisions needed"

- **First meeting with a new client:** Add “include rapport-building questions and discovery angle”

Playbook 2: Transform Raw Notes into Structured Minutes

When to use it:

You just finished a meeting and have messy, unstructured notes (or even a transcript). You need clean, professional meeting minutes within the hour.

The prompt:

```
Role: You are a professional meeting secretary.

Context: I just finished a [type] meeting with
[attendees/teams]. Below are my raw notes.

Transform them into structured meeting minutes:
1. Meeting header (date, attendees, objective)
2. Key Decisions (numbered, with rationale)
3. Action Items (table: action, owner, deadline)
4. Open Questions (unresolved items needing follow-up)
5. Key Takeaways (3-5 sentences for someone who
   wasn't there)

Rules:
- If a decision is ambiguous in my notes, flag it
  as "[TO CONFIRM]" rather than guessing.
- Keep the tone factual and concise.
- Don't invent information that isn't in my notes.

[PASTE RAW NOTES HERE]
```

What you get:

Professional meeting minutes you can send directly to attendees. Action items are in a table format, decisions are numbered and unambiguous, and anything unclear is flagged for your review.

Variants:

- **Workshop or brainstorm:** Add "group ideas by theme and rank by feasibility"

- **Board or steering committee:** Add "include a formal resolution section and voting outcomes"
- **Informal call with quick follow-ups:** Simplify to "just extract action items and next steps, skip the formal structure"

> **From the author**
> I personally have this playbook built into a dedicated skill. Every time I need meeting minutes, I paste my raw notes and type /meeting-notes — that's it. The skill handles the RCFC structure, applies my formatting preferences, and delivers clean minutes in seconds. If you find yourself using a playbook more than twice a week, that's your signal to turn it into a skill.

Playbook 3: Draft a Difficult Email

When to use it:

You need to write an email that requires careful wording: a polite decline, a firm push-back, a bad-news delivery, a salary negotiation, or a sensitive follow-up.

The prompt:

```
Role: You are a senior communications advisor with
expertise in [professional/diplomatic] correspondence.

Context:
- Situation: [describe what happened and why you need
  to write this email]
- Recipient: [who they are, your relationship, their
  likely emotional state]
- My goal: [what I want to achieve with this email]
- Tone: [firm but respectful / empathetic but clear /
  diplomatic / direct]

Write the email. Include a subject line.

Constraints:
- Length: 100-200 words (respect their time)
- Don’t apologize excessively
- Don’t make promises I can’t keep
- End with a clear next step
- [Any other specific constraints]
```

What you get:

A polished, send-ready email that balances your goal with the relationship dynamics. Subject line included, next step clear, tone calibrated to the situation.

Variants:

- **Declining a meeting or request:** Add “offer a brief reason without over-explaining, and suggest an alternative”

- **Delivering bad news to a client:** Add “acknowledge the impact on their side before explaining the situation”
- **Following up after no response:** Add “keep it light, assume good intent, make it easy for them to reply”
- **Negotiating terms or pricing:** Add “anchor on value delivered, not on cost, and propose two options”

Playbook 4: Analyze a Long Document and Extract the Essentials

When to use it:

You have a 20-page report, a contract, a policy document, or a lengthy email thread, and you need the key takeaways fast.

The prompt:

```
Role: You are a senior analyst with expertise in
[relevant domain].

I'm going to share a [document type]. I need you to:

1. Provide a 3-sentence executive summary
2. Extract the [5] most important findings or clauses
3. Identify any risks, contradictions, or red flags
4. List any action items or decisions required from me
5. Highlight anything that seems unusual compared to
   standard [industry] practice

Rank findings by [impact / urgency / relevance to
my role].

Be direct. One paragraph per finding. If something
is ambiguous, say so rather than guessing.

[PASTE OR UPLOAD DOCUMENT]
```

What you get:

A structured analysis that turns a 2-hour reading task into a 5-minute review. Key findings are ranked, red flags are surfaced, and action items are explicit.

Variants:

- **Contract review:** Add "flag non-standard clauses, liability caps, termination conditions, and IP ownership"
- **Financial report:** Add "cross-reference management's narrative with the actual numbers"

- **Competitor analysis:** Add “compare against our positioning and flag strategic implications”

Playbook 5: Create a Presentation Outline and Content

When to use it:

You need to build a slide deck and want the AI to generate the full structure and content before you open PowerPoint, Google Slides, or Canva.

The prompt:

```
Role: You are a presentation strategist who has
designed decks for [consulting firms / executives /
sales teams].

Context:
- Topic: [what the presentation is about]
- Audience: [who will see it and what they care about]
- Purpose: [inform / persuade / get approval / sell]
- Duration: [5 / 15 / 30 minutes]

Create a slide-by-slide outline with:
- Slide title
- 3-4 bullet points of content
- Speaker notes (50-75 words, what I should say)
- Suggested visual (chart type, diagram, image idea)

Structure the narrative with a clear hook, build-up,
and call to action.

Constraints:
- Maximum [10/15/20] slides
- No more than 6 words per title
- No more than 4 bullets per slide
- [Any brand or style constraints]
```

What you get:

A complete presentation blueprint you can feed directly into your preferred slide tool. Each slide has content, notes, and visual suggestions. The narrative arc is built in.

Variants:

- **Investor pitch deck:** Add "follow the standard VC structure: problem, solution, market, traction, team, ask"
- **Internal training:** Add "include knowledge check questions every 4-5 slides"
- **Client proposal:** Add "lead with their pain points, position our solution, end with ROI and next steps"

Playbook 6: Run a Weekly Industry Watch

When to use it:

You want to stay current on your industry without spending hours reading newsletters and articles every week.

The prompt:

```
Role: You are a senior industry analyst specializing
in [your industry/domain].

Search the web for the most important developments
in [your industry] from the past 7 days.

Create a weekly briefing with:
1. Top 5 developments (one paragraph each, ranked by
   impact on [my specific role/company]). For each
   development, include the source URL as a clickable
   link so I can read the full article.
2. One emerging trend to watch (with source link)
3. One competitor move worth noting (with source link)
4. Implications for my work: what should I be thinking
   about or acting on this week?

Sources: prioritize [industry publications, major
business press, regulatory announcements].
Tone: analytical, not sensational.
Length: under 500 words total.
```

What you get:

A concise weekly briefing that keeps you informed in 5 minutes. Findings are ranked by relevance to your role, not by headline appeal.

Variants:

- **Technology watch:** Change focus to "new tools, frameworks, and product launches relevant to my stack"

- **Regulatory watch:** Add “focus on regulatory changes, compliance updates, and policy proposals”
- **Competitive intelligence:** Add “track specific competitors: [names], focusing on hiring, product launches, and partnerships”

Playbook 7: Prepare for an Interview (Giving or Receiving)

When to use it:

You have an interview coming up — either you're interviewing a candidate, or you're the candidate. You need targeted preparation.

The prompt:

```
Role: You are a senior HR advisor with 15 years of
experience in [industry] recruitment.

Context: I am [interviewing a candidate for X role /
preparing for an interview at X company for Y role].

[IF INTERVIEWER]
- Job description: [paste or describe]
- Create 10 interview questions: 5 technical,
  3 behavioral (STAR format), 2 culture fit
- For each: what a good answer looks like vs.
  a red flag answer
- Suggest a 45-minute interview structure

[IF CANDIDATE]
- Job description: [paste or describe]
- My background: [key highlights or paste CV]
- Generate 10 likely questions with suggested answers
  tailored to my experience
- Identify 3 areas where I might be challenged
  and how to address them
- Suggest 5 smart questions for me to ask them
```

What you get:

A complete interview prep kit. If you're the interviewer: structured questions with evaluation criteria. If you're the candidate: predicted questions with tailored answers and a list of smart questions to ask.

Variants:

- **Technical role:** Add “include a practical exercise or case study suggestion”
- **Leadership role:** Add “focus on strategic thinking, stakeholder management, and change leadership”
- **Panel interview:** Add “antipate different perspectives from different panelists (HR vs. technical vs. business)”

Playbook 8: Debug a Process or Diagnose a Problem

When to use it:

Something isn't working as expected in your business process, project, or system. You need a structured root-cause analysis.

The prompt:

```
Role: You are a senior operations consultant
specializing in process improvement and root-cause
analysis.

Context:
- Process/system: [describe what should be happening]
- Problem: [describe what is actually happening]
- Since when: [when did the problem start]
- What changed: [any recent changes to the process,
  team, tools, or inputs]
- Impact: [how this is affecting the business]

Perform a structured root-cause analysis:
1. Restate the problem clearly
2. Generate 5 possible root causes (use 5 Whys or
   Ishikawa approach)
3. For each cause: how to verify it (what data to
   check, who to ask)
4. Rank causes by likelihood and impact
5. Recommend immediate actions and long-term fixes

Be methodical. Don't jump to solutions before
analyzing causes.
```

What you get:

A structured diagnostic that moves you from "something's broken" to "here are the three most likely causes and how to verify each one." Clear next steps for both short-term fixes and permanent solutions.

Variants:

- **Software bug:** Change role to “senior software engineer” and add “include log analysis approach and debugging steps”
- **Team performance issue:** Change role to “experienced team lead” and add “consider people factors: workload, clarity, motivation, skills gap”
- **Financial discrepancy:** Change role to “senior auditor” and add “focus on reconciliation steps, data lineage, and control gaps”

Playbook 9: Write a Proposal or Quoted Offer

When to use it:

You need to send a client a proposal, a quotation, or a scoped offer for a project or engagement.

The prompt:

```
Role: You are a senior business development consultant
with expertise in writing winning proposals.

Context:
- Client: [description]
- What they asked for: [their need or RFP summary]
- What I'm proposing: [my solution, approach, scope]
- My differentiators: [why me vs. competitors]
- Budget range: [if known]
- Timeline: [proposed timeline]

Write a professional proposal with these sections:
1. Executive Summary (their problem, our solution,
   expected outcome — under 150 words)
2. Understanding of the Need (show I understand
   their situation)
3. Proposed Approach (methodology, phases, milestones)
4. Deliverables and Timeline (table format)
5. Investment (pricing, payment terms, what's included)
6. Why Us (proof points, relevant experience, team)
7. Next Steps (clear call to action)

Constraints:
- Total length: 3-5 pages
- Tone: confident, solution-focused, not salesy
- Lead with value, not features
- No jargon the client wouldn't understand
```

What you get:

A polished proposal you can put on company letterhead and send. The structure leads the client from “you have a problem” to “here’s how we solve it” to “here’s what it costs” to “here’s what to do next.”

Variants:

- **Freelance SOW:** Simplify to “2-page scope of work with deliverables table, rate, and terms”
- **RFP response:** Add “match the RFP structure exactly and reference their evaluation criteria”
- **Renewal or upsell:** Add “lead with results from the current engagement and position the expansion as a natural next step”

Playbook 10: Build an Action Plan from a Vague Situation

When to use it:

You're facing a complex or ambiguous situation and you need to move from "I don't know where to start" to "here's my plan."

The prompt:

```
Role: You are a senior strategy consultant who
specializes in turning ambiguous situations into
actionable plans.

Context:
- Situation: [describe the situation as honestly as
  you can, including what's unclear]
- Goal: [what does success look like, even roughly]
- Constraints: [time, budget, resources, dependencies]
- What I've already tried or considered: [if anything]

Help me build a structured action plan:
1. Clarify the situation: restate the problem in one
   sentence. If there are multiple problems, separate
   them.
2. Identify the 3 most critical unknowns (things I
   need to figure out before I can act)
3. For each unknown: suggest how to resolve it
   (who to talk to, what data to gather, what to test)
4. Propose a phased action plan:
   - This week: immediate actions
   - Next 2 weeks: foundational steps
   - Next month: execution milestones
5. Identify the biggest risk and a mitigation strategy

Be honest. If something in my situation doesn't
make sense, say so.
```

What you get:

A structured action plan that breaks an overwhelming situation into concrete, time-bound steps. Critical unknowns are identified and paired with resolution strategies. You move from paralysis to momentum.

Variants:

- **Career decision:** Add "use a decision matrix with weighted criteria and help me score each option"
- **Project rescue:** Add "focus on what's salvageable, what to cut, and how to reset stakeholder expectations"
- **New business initiative:** Add "use a lean canvas approach: problem, solution, key metrics, unfair advantage"

HOW TO GET THE MOST FROM THESE PLAYBOOKS

These ten playbooks are starting points, not finished products. Here's how to make them truly yours:

- **Adapt the role** to match your specific domain. "Senior analyst" is fine, but "Senior supply chain analyst with manufacturing experience" is better.
- **Add your context files.** When running these inside a Project with your Identity Card and Style Guide loaded, the outputs will already match your voice and vocabulary.
- **Save the ones you use most as skills.** If you run Playbook 2 (meeting minutes) three times a week, turn it into a dedicated skill so you can trigger it with a single phrase.
- **Iterate and refine.** After each use, note what worked and what didn't in your Preferences Journal (File 5 from Part 3). The prompts get better every time.

> **The compound effect**
> A professional who uses even 3-4 of these playbooks consistently saves 5-8 hours per week. Over a month, that's an extra full work week. Over a year, it's two months of recovered productive time. The ROI of this guide pays for itself after one use of one playbook.

PART 6
Level Up: From User to Architect

THE FEEDBACK LOOP: HOW YOUR AI GETS BETTER OVER TIME

If you've followed this guide from Part 0 to Part 5, you're already ahead of 90% of professionals using AI. You have context files, a prompt method, playbooks, and you know how to produce visual deliverables. But there's a difference between using AI and having a system that improves itself.

The professionals who extract the most value from AI don't just use it — they run a continuous feedback loop:

1. **Prompt** — you send a request using your playbook or skill.
2. **Output** — the AI delivers a result.
3. **Evaluate** — you assess: what was good? What missed the mark? Was the tone right? Was anything missing?
4. **Refine** — you tell the AI what to fix, or you adjust your prompt for next time.
5. **Record** — you update your Preferences Journal (File 5) with what you learned. "When writing client emails, always open with the key decision, not the background."

Step 5 is what separates a user from an architect. Most people do steps 1-4 and then forget about it. The ones who record their learnings build a compounding advantage: every week, their prompts get better, their context files get sharper, and their outputs require less editing.

> **The Friday 5-minute habit**
> Every Friday before closing your laptop, spend 5 minutes reviewing your AI conversations from the week. Ask yourself: Did I discover a prompt pattern that worked well? Did the AI make a recurring mistake I should add to my constraints? Is there a task I repeated three times that should become a skill? Update your Preferences Journal and your context files accordingly. Five minutes a week. The compound effect over a year is enormous.

THE THREE LEVELS OF AI MATURITY

Not everyone needs to reach Level 3. But knowing where you are — and where you could go — helps you decide where to invest your time.

Level 1: Occasional User

You're here if: you use AI a few times a week for isolated tasks. You open a chat, get an answer, close the tab.

What it looks like:

- You use 1-2 playbooks from Part 5 regularly.
- You have Custom Instructions or a basic context set up.
- You save maybe 2-3 hours per week.

How to stay effective here:

Keep your Preferences Journal updated. Use the RCFC method for anything beyond a quick question. That alone puts you ahead of most people.

Time saved: 2-4 hours/week.

Level 2: Integrated Workflow

You're here if: AI is a systematic step in your daily work. You don't think "should I use AI for this?" — it's just part of how you work.

What it looks like:

- You have multiple Projects/GPTs set up for different workstreams.
- You've built 3-5 custom skills for your recurring tasks.
- You've connected 2-3 external tools via connectors (Drive, email, Notion).

- Your context files are mature and regularly updated.
- You save 5-8 hours per week consistently.

How to get here from Level 1:

1. Identify your 5 most time-consuming recurring tasks.
2. Turn the top 3 into dedicated skills or playbooks.
3. Set up a Project for each major workstream (client work, internal, business development).
4. Connect at least one external tool (Google Drive is the highest-ROI starting point).
5. Run the Friday 5-minute review for a month.

Time saved: 5-10 hours/week.

Level 3: Autonomous Systems

You're here if: you have AI workflows that run without your direct involvement. Things happen in the background while you focus on high-value work.

What it looks like:

- You have automated pipelines: email digests that generate themselves, reports that compile overnight, monitoring that alerts you when something needs attention.
- You use workflow tools (n8n, Make, Zapier) connected to AI APIs.
- You've moved from "I use AI" to "AI runs parts of my business."
- You save 10+ hours per week — and some tasks simply didn't exist before because they weren't feasible manually.

What it takes to get here:

Level 3 requires some technical comfort. You don't need to be a developer, but you need to be willing to learn basic automation tools. The concepts are accessible — modern no-code platforms have made this dramatically easier — but it's a genuine investment of time (10-20 hours to set up your first automated workflow).

Examples of Level 3 workflows:

- Every Sunday at 10am, an automated pipeline scans your subscribed newsletters, extracts the top insights, and saves a structured weekly briefing to your Notion workspace.
- When a new client inquiry arrives in your inbox, an agent extracts the key details, searches your CRM for existing contacts, and drafts a personalized response for your review.
- A monthly financial report is automatically generated: data pulled from your spreadsheet, analysis run by AI, formatted document delivered to your inbox.

Time saved: 10-20+ hours/week, plus new capabilities that weren't possible before.

The Maturity Progression at a Glance

	LEVEL 1	LEVEL 2	LEVEL 3
Label	Occasional User	Integrated Workflow	Autonomous Systems
AI usage	A few times/week	Multiple times/day	Runs in background
Skills	0-1	3-5	10+
Connectors	0	2-3	5+
Automation	None	Manual triggers	Scheduled pipelines
Time saved	2-4 hrs/week	5-10 hrs/week	10-20+ hrs/week
Setup effort	1-2 hours	1-2 days	1-2 weeks
Technical level	None	Low	Moderate

> **Where should you aim?**
> For most professionals, Level 2 is the sweet spot. It delivers massive time savings with modest setup effort and no technical skills required. Level 3 is for those who want to push further and are willing to invest the time — or who have tasks that truly demand automation (high-volume, repetitive, data-heavy). Don't rush to Level 3. Get comfortable at Level 2 first. The jump will feel natural when you're ready.

WHEN TO INVEST IN A PAID SUBSCRIPTION (AND WHICH ONE)

All three major platforms offer free tiers. The question isn't whether you can afford to pay — it's whether the free tier is costing you more in limitations than the subscription costs in money.

Signs You've Outgrown the Free Tier

- You hit message limits multiple times per week and have to wait or switch tools.
- You need to upload files regularly and the free tier restricts it.
- You want to use Projects/Custom GPTs/Gems but they're locked behind the paywall.
- You need the latest model for complex reasoning tasks and only get the basic model for free.
- You're using AI for client-facing work and the quality difference between free and paid models is visible.

The Decision Matrix

	CHATGPT PLUS	CLAUDE PRO	GEMINI ADV.
Price	$20/month	$20/month	$20/month
Best model	GPT-5.5	Claude Opus 4.7	Gemini 3.1 Pro
Projects/GPTs	Yes	Yes (Projects)	Yes (Gems)
File uploads	Generous	Generous	Generous
Image generation	GPT Image	No	Nano Banana
Presentation files	Agent mode .pptx	Native .pptx	Google Slides
Connectors	Plugins + Apps	MCP connectors	Google Workspace
Code execution	Yes (Python)	Yes (artifacts)	Yes (Python)
Best value if...	You need images + versatility	You need writing + analysis	You're in Google ecosystem

> **The math is simple**
> If you save 5 hours per week using AI, and your time is worth more than $4/hour, a $20/month subscription pays for itself in the first week. For most professionals, the paid tier isn't an expense — it's one of the highest-ROI investments you can make in your own productivity.

5 RESOURCES TO KEEP LEARNING

The AI landscape moves fast. Here are five high-quality, low-noise sources to stay current without drowning in hype.

1. Anthropic's Prompt Engineering Guide

The official guide from the makers of Claude. Covers everything from basic prompting to advanced techniques. It's the most rigorous, no-nonsense resource available.

docs.anthropic.com/en/docs/build-with-claude/prompt-engineering

2. OpenAI's Cookbook

A collection of practical examples and best practices from the ChatGPT team. Particularly useful for code-related tasks and API integrations.

cookbook.openai.com

3. r/ChatGPT and r/ClaudeAI on Reddit

The most active communities of real users sharing real use cases. Filter by "Top — This Month" to find the best content without wading through noise. Great for discovering workflows you hadn't thought of.

reddit.com/r/ChatGPT • reddit.com/r/ClaudeAI

4. Ethan Mollick's Newsletter ("One Useful Thing")

A Wharton professor who writes about AI from a practical, business-oriented perspective. No hype, no fear-mongering, just thoughtful analysis of what AI means for the way we work.

oneusefulthing.org

5. The Official Documentation of Your Chosen Platform

Seriously. The Claude docs, the ChatGPT help center, and the Gemini guides are better than 90% of the tutorials out there. When in doubt, go to the source. Features change every month — the official docs are always up to date.

docs.claude.com • help.openai.com • gemini.google.com/help

What to ignore
Skip the "10x your productivity with this ONE prompt" threads. Skip the AI influencers selling courses on how to sell AI courses. Skip anything that promises magic without effort. The professionals getting real results are the ones who built a system, practiced it consistently, and refined it over time. That's exactly what you're doing right now.

What's next: your personalized setup
You now have the knowledge, the method, and the tools. Part 7 brings it all together: a dedicated skill that analyzes your profession and builds your complete AI architecture — projects, skills, connectors, and reference files, all tailored to the way you work. It's the final step from "user" to "architect."

PART 7

The Perfect Setup for You

Including the "Setup My Own LLM" Skill

YOUR AI IS NOT A TOOL. IT'S A WORKSPACE.

This is the final chapter, and it's where everything converges. You've learned how to prompt (Part 2), how to build persistent context (Part 3), how to produce deliverables (Part 4), and you have playbooks for your most common tasks (Part 5). Now it's time to architect the whole thing into a coherent system.

Think about your physical workspace. You have a desk organized by project. You have folders for different clients. You have templates for recurring documents. You have tools within arm's reach. You didn't build this setup in a day — it evolved over months. But once it was in place, it made you dramatically faster.

Your AI workspace works the same way. Instead of one generic chatbot, you're going to build a structured environment with dedicated projects, custom skills, connected tools, and reference files — all tailored to the way you actually work.

The Anatomy of a Professional AI Setup

A complete setup has five layers. Each one builds on the previous:

LAYER	WHAT IT IS	EXAMPLE	SET UP IN PART
1. Instructions	Global preferences that apply everywhere	Your name, language, date format, tone	Part 1 (Mode 2)
2. Projects	Dedicated workspaces for different areas	"Client Work," "Business Dev," "Admin"	Part 1 (Mode 3)
3. Context files	Reference docs loaded into each project	Identity Card, Style Guide, Glossary	Part 3
4. Skills	Reusable capabilities triggered on demand	RCFC Builder, Meeting Notes, Proposal Drafter	Parts 2, 3, 7
5. Connectors	Links to your external tools	Google Drive, Gmail, Notion, Slack	Part 1 (Mode 5)

The skill you're about to install will design layers 2 through 5 for you, based on your specific profession.

HOW TO ORGANIZE YOUR WORK INTO PROJECTS

The biggest mistake people make is throwing everything into one workspace. That's like having one folder on your desktop called "Stuff." It works for a week. After that, it's chaos.

The right way to split your work into projects depends on your role. Here are three common patterns:

Pattern A: Split by Client or Stakeholder

Best for: consultants, freelancers, agency professionals, account managers.

Each client gets their own project. The project contains that client's context (industry, key contacts, past deliverables, communication preferences), plus the skills and templates specific to that engagement.

```
Project: Acme Corp
  ├─ Context: client profile, key contacts, past proposals
  ├─ Skills: proposal drafter, status report generator
  └─ Connectors: Google Drive (Acme folder)

Project: GlobalTech
  ├─ Context: client profile, technical glossary, SOW
  ├─ Skills: meeting notes, risk tracker
  └─ Connectors: Google Drive (GlobalTech folder)
```

Pattern B: Split by Workflow

Best for: managers, specialists, and anyone with distinct work modes.

Each major type of work gets its own project. The project is optimized for that workflow's specific needs.

```
Project: Client Delivery
  ├─ Context: delivery standards, SLA templates
  ├─ Skills: status report, deliverable reviewer
  └─ Connectors: project management tool
```

```
Project: Business Development
  ├─ Context: value prop, pricing, case studies
  ├─ Skills: proposal drafter, pitch preparer
  └─ Connectors: CRM, email

Project: Internal / Admin
  ├─ Context: company policies, org chart
  ├─ Skills: meeting notes, email drafter
  └─ Connectors: calendar, Slack
```

Pattern C: Hybrid (Client + Workflow)

Best for: professionals who juggle multiple clients AND have distinct internal workflows.

Major clients get their own project. Everything else is organized by workflow. This is the most common pattern for experienced freelancers and senior consultants.

> **Start with 3 projects. No more.**
> Don't over-architect on day one. Create three projects that match your most frequent work modes. Use them for two weeks. Then add more as needed. The Setup My Own LLM skill will recommend a starting structure, but you always have the final say.

WHICH SKILLS TO BUILD FOR YOUR PROFESSION

A skill should exist for any task you do at least twice a week that follows a repeatable pattern. The test is simple: if you find yourself re-typing the same kind of prompt with minor variations, that's a skill waiting to be created.

How to Identify Your Top Skill Candidates

1. List the 10 tasks you do most often in a typical week.
2. For each, ask: does this follow a pattern? Could I describe the steps to someone?
3. For each "yes," ask: how much time does it take me today? How much could AI save?
4. Rank by (frequency × time saved). The top 3-5 are your first skills.

Skill Ideas by Profession

Here are concrete skill ideas for common professional profiles. The Setup My Own LLM skill will generate personalized recommendations based on your actual work.

Consultant / Advisory

- Client Meeting Prep — generates briefing sheets from meeting context
- Proposal Drafter — takes client need + your approach, outputs a structured proposal
- Deliverable Reviewer — checks a document against your quality standards
- Status Report Generator — turns bullet points into formatted status updates
- Invoice Narrator — transforms timesheet data into a justified invoice summary

Manager / Team Lead

- 1:1 Prep — generates discussion points based on team updates and goals
- Decision Memo Writer — structures a decision with options, criteria, and recommendation
- Team Update Compiler — aggregates inputs from multiple people into one summary

- Feedback Drafter — writes constructive feedback based on situation and desired outcome

Finance / Controlling

- Variance Commentator — explains budget vs. actual deviations in narrative form
- Board Deck Preparer — structures financial data into presentation-ready content
- Policy Summarizer — extracts key changes from updated regulations or standards

Creative / Marketing

- Content Calendar Builder — generates a month of post ideas from themes and goals
- Brief Translator — turns a vague client brief into an actionable creative brief
- Copy Adapter — rewrites content for different channels (LinkedIn, email, blog)

> **The skill anatomy reminder**
> Every skill needs: a name, a description, trigger phrases (the words that activate it), and the full instructions (the prompt the AI follows when triggered). The Setup My Own LLM skill generates all of this for you, including a ready-to-install SKILL.md file for each recommended skill.

WHICH CONNECTORS TO SET UP FIRST

Connectors (MCP integrations) are the bridge between your AI and your real tools. Start with the one that removes the most friction.

The Priority Order

#	CONNECTOR	WHAT IT UNLOCKS	CONNECT IF…
1	Google Drive / OneDrive	"Find and summarize the Q3 report"	You store files in the cloud
2	Gmail / Outlook	"Draft a reply to Sarah's last email"	You spend >30 min/day on email
3	Calendar	"Prep me for tomorrow's meetings"	You have 3+ meetings per day
4	Notion / project tools	"Update the status of task X"	You manage tasks or docs in Notion
5	Slack / Teams	"Summarize the #project channel"	You're in 5+ active channels

> **From the author: a word on emails**
> Most people are eager to use AI to answer emails. I personally don't. In my experience, the AI is often missing the mark — either too generic, too polished, or simply unaware of the real dynamics at play in the conversation. It doesn't know the history with that person, the political context, or what I actually want to signal between the lines.
>
> Where I do use AI for emails is as a format manager, not a content writer. I draft the email myself — because I know what I want to say — and then ask the AI to refine it: make it more polite, fix typos, tighten the structure, adjust the tone for the recipient. That's where it genuinely helps without the risk of sending something that doesn't sound like me.
>
> *Bottom line: connect your email if you want, but think of AI as your editor, not your ghostwriter.*

Start with connector #1. Use it for a week. Then add the next one.

YOUR THIRD SKILL: SETUP MY OWN LLM

This is the most ambitious skill in the guide. It's an AI architect that interviews you about your profession, analyzes your work patterns, and designs a complete, personalized AI environment.

Think of it as hiring a consultant to set up your digital workspace — except the consultant is the AI itself, and the setup takes 30 minutes instead of 30 days.

How It Works

1. **Discovery** — the skill asks 10-15 questions about your work: role, clients, recurring tasks, tools, pain points.
2. **Architecture** — proposes a project structure: named projects, each with a clear scope.
3. **Content** — for each project, lists context files, connectors, which Part 3 files to attach, and writes tailored project instructions using the RCFC method (or whichever prompting structure best fits the project's purpose) so the AI knows how to behave the moment you open a conversation.
4. **Skills Design** — proposes 3-5 custom skills with complete SKILL.md skeletons.

The Skill File: Setup My Own LLM

Copy the following text. Paste it into a Claude conversation and ask Claude to turn it into a downloadable skill file. Then go to Customize → Skills → Create skill → Upload skill.

```
---
name: setup-my-own-llm
description: >
  Interviews the user about their professional context,
  then designs a complete AI workspace: project structure,
  context files per project, recommended skills with
  SKILL.md skeletons, and connector strategy. Adapts to
  any profession.
---
```

```
# Setup My Own LLM

## Purpose
Design a complete, personalized AI workspace based on
the user's profession, work patterns, and tools. Output
a ready-to-implement architecture: projects, context
files, skills, and connectors.

## Trigger phrases
Activate when the user says any of:
- "setup my LLM"
- "design my AI workspace"
- "build my AI setup"
- "architect my AI environment"
- "configure my AI"
- "create my perfect setup"

## Workflow

### Phase 1: Discovery (10-15 questions in 3 batches)

Start with one opening question:
"Describe your professional role in 2-3 sentences:
what you do, who you serve, and whether you work
solo, in a team, or as a freelancer."

Based on the answer, classify the user as one of:
- CONSULTANT (external advisor, client-facing)
- MANAGER (team lead, department head)
- FREELANCER (independent operator)
- SPECIALIST (individual contributor, deep expertise)
- CREATIVE (content, marketing, design)

Then ask 3 batches. Wait between each.

**Batch 1 - Work Structure (4-5 questions):**
- How many active clients/projects do you manage
  at any given time?
```

```
- What are the 3 biggest categories of work in
  your typical week?
- What tools do you use daily? (email, calendar,
  project management, CRM, storage, communication)
- Do you work primarily in English, or do you
  need multilingual support?

**Batch 2 - Tasks and Pain Points (4-5 questions):**
- What are the 5 tasks you spend the most time
  on each week?
- For each: roughly how long does it take, and
  does it follow a repeatable pattern?
- Which 3 tasks frustrate you the most?
- What documents or deliverables do you produce
  most often?

**Batch 3 - Tools and Ambitions (3-4 questions):**
- Which connectors would be most valuable for you?
  (Google Drive, Gmail, Calendar, Notion, Slack, other)
- Do you already have context files set up, or are
  we starting from scratch?
- On a scale of 1-5, how comfortable are you with
  technology and AI tools?
- What would "success" look like for your AI setup
  in 30 days?

### Phase 2: Architecture Proposal

Based on all answers, generate:

**A. Project Structure**
Propose 3-5 projects. For each:
- Project name (clear, action-oriented)
- Scope description (one sentence)
- Which context files to include
- Which connectors to activate
- Pattern used: by client, by workflow, or hybrid
- **Project instructions**: write a ready-to-paste
  instruction block for each project. Use the RCFC
  method or any prompting structure that best fits
```

```
  the project's purpose. These instructions tell
  the AI how to behave in every conversation within
  that project (default role, tone, output format,
  standing constraints). The user should be able to
  copy-paste these directly into the project's
  "Custom instructions" field.

Present in a clear table format.

**B. Shared vs. Project-Specific Files**
- Shared files (loaded in ALL projects): typically
  Identity Card, Style Guide, Preferences Journal
- Project-specific files: typically Glossary and
  Templates (adapted per client/domain)

Ask: "Does this structure match how you think about
your work? Want me to adjust anything before I
design the skills?"

### Phase 3: Skills Design

Propose 3-5 custom skills. For each, generate:

1. **Skill name** (kebab-case, descriptive)
2. **Description** (what it does and when to use it)
3. **Trigger phrases** (3-5 phrases)
4. **Which project(s) it belongs to**
5. **SKILL.md skeleton** (ready to install):
   - Purpose section
   - Trigger phrases section
   - Workflow section (phases and steps)
   - Question bank (if the skill asks questions)
   - Rules section

The SKILL.md should be complete enough to install
and use immediately. Generate real, actionable
instructions based on the user's stated work.

### Phase 4: Connector Roadmap
```

```
Recommend a connector activation sequence:
- Week 1: connect [highest-impact connector]
- Week 2: connect [second connector]
- Month 1: connect [remaining connectors]

For each, explain what it unlocks and how to set
it up (platform-specific steps).

### Phase 5: Implementation Checklist

Generate a numbered checklist:
1. Create projects [list names]
2. Upload shared context files to all projects
3. Create project-specific context files
4. Install skills in their respective projects
5. Connect [connector 1]
6. Test: run [specific playbook] in [project name]
7. Iterate: use for one week, then review

## Rules
- Never propose more than 5 projects initially.
- Never propose more than 5 skills initially.
- Always explain WHY a project or skill is
  recommended, linking to stated pain points.
- Generate REAL skill content, not placeholders.
- Use the same language as the user.
- If the user already has context files, reference
  them by name and suggest which project they go to.
- End with: "Your setup is ready. Start with
  [project name] and [skill name] today. The rest
  can wait until next week."
```

Installation Guide

For Claude:

- Copy the skill text above and paste it into a Claude conversation with: “Turn this into a proper skill file I can download.”
- Download the generated file.
- Go to Customize → Skills → Create skill → Upload skill.
- Start a conversation and type “setup my LLM.”
- Answer the three batches of questions.
- Review the proposed architecture, then follow the implementation checklist.

For ChatGPT:

- Create a Custom GPT with the skill instructions.
- Run the setup conversation.
- Implement the output by creating separate Custom GPTs for each proposed project.

For Gemini:

- Create a Gem with the skill instructions.
- Run the setup conversation.
- Create separate Gems for each proposed project and paste the relevant context.

EXAMPLE: WHAT THE SKILL PRODUCES

Here's a condensed example for a freelance management consultant.

User Profile (from Discovery)

```
Role: Freelance operations consultant, 12 years exp.
Clients: 2-3 active at any time, mid-size manufacturers
Top tasks: client reports, meeting prep, proposals,
  email, process documentation
Pain points: re-explaining context, writing proposals
  from scratch, formatting status reports
Tools: Gmail, Google Drive, Google Calendar, Notion
Tech comfort: 3/5
```

Proposed Architecture

Projects:

PROJECT	SCOPE	CONNECTORS
Client: Acme Corp	All work for Acme — reports, meetings, docs	Drive (Acme folder), Gmail
Client: TechFlow	All work for TechFlow — same structure	Drive (TechFlow folder), Gmail
Business Dev	Proposals, pitches, lead follow-ups	Gmail, Calendar
Admin & Finance	Invoicing, contracts, personal admin	Gmail, Calendar

Shared files across all projects: Identity Card, Style Guide, Preferences Journal.

Project-specific: each client project gets its own Glossary and Templates.

Recommended Skills:

SKILL	WHAT IT DOES	PROJECT(S)

client-meeting-prep	Generates a briefing sheet from meeting invite + past notes	Client projects
status-report-builder	Turns bullet points into formatted weekly reports	Client projects
proposal-drafter	Interviews for client need, outputs a structured proposal	Business Dev
email-tone-adapter	Rewrites drafts with the right tone for the recipient	All projects
process-documenter	Structures raw process observations into clean SOP format	Client projects

Each skill comes with a complete SKILL.md file ready to install — generated in full by the skill during the actual conversation.

Implementation Checklist:

1. Create the 4 projects in Claude.
2. Upload Identity Card, Style Guide, and Preferences Journal to all 4.
3. Create client-specific Glossary and Templates for Acme and TechFlow.
4. Install the 5 skills in their respective projects.
5. Connect Google Drive — the highest-impact connector for your workflow.
6. Test: open the Acme project, paste your latest meeting notes, and type /meeting-prep.
7. Use for one week. Update your Preferences Journal on Friday.

YOU'RE AN ARCHITECT NOW

If you've followed this guide from Part 0 to here, you've done something that fewer than 5% of AI users ever do: you've built a system.

You didn't just learn how to write prompts. You built a structured environment where the AI knows your role, speaks your language, follows your standards, and handles your recurring tasks — all before you type a single word.

That's the difference between someone who uses AI and someone who has an AI copilot.

What You've Built

- A clear understanding of what AI does well and where it falls short (Part 1).
- A repeatable method for writing professional prompts — RCFC (Part 2).
- A persistent context system that eliminates cold starts forever (Part 3).
- The ability to produce visual deliverables directly from the chat (Part 4).
- 10 ready-to-use playbooks for your most common tasks (Part 5).
- A roadmap from occasional user to integrated workflow to autonomous systems (Part 6).
- A fully personalized AI architecture: projects, skills, connectors, and reference files (Part 7).

What to Do Next

1. **Run the Setup My Own LLM skill today.** Implement its recommendations this week.
2. **Start with one project and one skill.** Don't try to build everything at once.
3. **Use the Friday 5-minute review** (Part 6) to continuously improve your setup.
4. **Share what works.** If a colleague or friend is struggling with AI adoption, share one playbook. One skill. One quick win.

> **One final thought**
> The technology will keep evolving. Models will get smarter. New features will launch. But the principles in this guide — structured prompting, persistent context, reusable skills, connected tools — are foundational. They apply to every AI platform, every model, and every upgrade cycle. You haven't just learned to use today's tools. You've learned to think in systems. That's the real skill, and it compounds forever.

BONUS CHAPTER

My Personal Setup

A look in the rearview mirror

The Professional's AI Playbook

WHY I'M SHARING THIS

Over the past two years, I've learned — through trial and error, dead ends, and small victories — everything you'll find in this guide. I've progressively shaped Claude into a genuine professional partner, adapted to my specific way of working.

Writing this guide gave me a reason to stop and take stock. When you're building something day after day, you rarely pause to look at the full picture. This chapter is that pause. It's a look in the rearview mirror — at the road I've traveled, the system I've built, and what I've learned along the way.

My hope is that by showing you a real, lived-in setup — not a theoretical one — you'll see what's possible when you commit to building a system over time. Not everything I did was efficient. Some things I'd do differently. But the result speaks for itself: AI went from a curiosity to the backbone of how I work.

MY JOB: FREELANCE SAP FICO CONSULTANT

I've been a freelance SAP FICO consultant for close to 10 years. If that means nothing to you, here's the plain-English version.

What is SAP?

SAP is the world's largest enterprise software company. Their flagship product — now called S/4HANA — is the operating system of large organizations: it manages finance, procurement, manufacturing, sales, logistics, and human resources in a single integrated platform. Think of it as the central nervous system of a company. When a factory produces a widget, SAP records the material consumed, calculates the cost, creates the financial postings, updates inventory, and triggers the invoice — all automatically.

What does FICO mean?

FICO stands for Financial Accounting (FI) and Controlling (CO) — the two core finance modules inside SAP. FI handles external reporting: general ledger, accounts payable, accounts receivable, asset accounting — everything the CFO needs to close the books and report to regulators. CO handles internal management accounting: cost centers, profit centers, product costing, profitability analysis — everything the COO and business unit leaders need to understand where money is made and lost.

My specialty within FICO is Product Costing, Material Ledger, and CO-PA (profitability analysis) — the part of SAP that tells a manufacturer exactly how much it costs to produce each product, and how profitable each product, customer, and market is.

The three roles I've held

Functional Expert

This is the deep specialist role. A functional expert is the person the team calls when there's a complex configuration problem, a tricky integration between modules, or a closing error nobody can figure out. You live in the SAP system itself — configuring business rules, troubleshooting issues, writing functional specifications for developers, and ensuring the system does what the business needs it to do. Think of it as the architect who also knows how to lay bricks.

Functional Integrator

The integrator role is broader. You're still technical, but your focus shifts to how different modules work together. When the sales team creates an order in SD and it needs to flow into finance in FI and impact the cost calculation in CO, someone needs to make sure all those connections work correctly. That's the integrator. You're the person who sees the big picture and makes sure the pieces fit — across modules, across teams, sometimes across countries.

Team Lead

The team lead manages a group of consultants (typically 3-8 people) on a project. You're still involved in the technical work, but you also run workshops with the client, coordinate deliverables, manage the timeline, escalate risks, and serve as the primary interface between the technical team and the client's leadership. It's where consulting meets management.

How AI fits into each role

In each of these roles, I've been able to put AI to work in different ways. Let me show you the before and after — because the contrast is where the value becomes obvious.

As a Functional Expert: from digging to asking

Before AI, troubleshooting an SAP error you'd never seen before looked like this:

- Step 1: Copy the error message and paste it into Google.
- Step 2: Scroll through results until you find a relevant SAP Community thread or an obscure blog post from 2014 written by a consultant in another country.
- Step 3: Try to decipher what the author was actually saying — half the time the explanation assumed a different system configuration, a different SAP release, or a different business context than yours.
- Step 4: Adapt the proposed solution to your own system and customizing. This is where things usually went wrong, because the blog post didn't mention the three prerequisites you needed to check first.

- Step 5: Test. Fail. Re-read the post. Try a different interpretation. Test again. Eventually get it right — and finally understand the root cause deeply, but only after losing 45 minutes to 2 hours.

Now with AI: I paste the SAP error message, my functional specification, and my unit test results. I ask for both the solution and a clear explanation of the underlying concept. Claude gives me a structured diagnostic: what's happening, why, what to check, and how to fix it — adapted to my system context because my project already knows my SAP configuration.

I go dramatically faster. And on top of the time saved, I avoid the headache that used to come with deciphering poorly written forum posts at 6pm on a Friday when the closing is blocked.

As an Integrator: catching up on your blind spots

The integrator role — sometimes called "Staff" in the product management world — puts you in an inherently uncomfortable position. You're responsible for making sure multiple modules work together, but you can't be a deep expert in all of them. Nobody is. The expertise in each SAP module goes so deep that mastering even one takes years. Mastering all of them is simply not realistic.

But here's the thing: to be a good integrator, you don't need full expertise in every module. You need a solid understanding of the landscape — what each module does, where the boundaries are, what data flows between them, and what's at stake when something changes. You need to get there quickly, and you need to ask the right questions.

This is where AI shines. It helps me catch up on domains outside my core expertise at a pace that was simply impossible before. When I need to understand how a change in MM pricing conditions cascades into FI account determination, I don't need to become an MM expert — I need the AI to teach me exactly what's relevant to my integration question. It fills my blind spots on demand, at the depth I need, when I need it.

This is exactly the kind of task I use the SAP Deep Dive Analyzer skill for. It maps the technical landscape so I can see the full picture without pretending to be an expert in areas I'm not.

As a Team Lead: escaping the formatting treadmill

In most team lead roles, a surprising amount of your time is spent moving the same information from one format to another. The content barely changes — but the container does, constantly.

Here's what a typical week looked like:

- Monday: present the project status in a PowerPoint during the steering committee.
- Tuesday: write a detailed follow-up email to the stakeholders, expanding on what was discussed during the presentation.
- Wednesday: extract parts of that email into a more concise update for a different audience.
- Thursday: gather new inputs from the team for the weekly status.
- Friday: compile everything into next Monday's PowerPoint, and the cycle starts again.

Before AI, this was all manual. Reformatting bullet points into prose, prose into slides, slides into emails. Low added-value work that you couldn't skip because the communication cadence demanded it. I used to spend 3-4 hours per week on pure formatting and reformatting.

Now with AI, I set up the formats I use (my slide template, my email structure, my status report layout). I explained the transformation workflow to the AI once: how information flows from the PPT to the email to the weekly update and back. Now, every week, with a few copy-pastes, I'm done. The AI handles the format conversion; I focus on the content and the decisions.

Bye-bye, lost hours on formatting. The Meeting Actions Tracker skill alone saves me at least 2 hours per week of post-meeting cleanup. Combined with the format transformation workflow, I've reclaimed nearly a full working day every week.

MY PROJECTS: HOW I ORGANIZED MY AI WORKSPACE

I use Claude as my primary platform, with a project architecture that follows the hybrid pattern described in Part 7 (split by workflow, with a naming convention to keep things organized). Each project is prefixed with "FL" for "freelance" — a small thing that helps when you're scanning a long list.

Here are my 7 projects, in the order I created them. Each one addresses a specific facet of my consulting work.

PROJECT	PURPOSE
FL01 — SAP FICO Technical Expert	My daily driver. This is where I go for any technical SAP question: table structures, transaction codes, configuration options, ABAP debugging hints, account determination logic. It's loaded with my SAP context files and connected to web search for SAP Notes lookup. I use it multiple times a day.
FL02 — S/4HANA Migration Advisor	Dedicated to ECC-to-S/4HANA migration topics. Simplifications, deprecated features, new Fiori apps, migration impact analysis. When a client asks "what changes if we move to S/4?" this is where I find the answer.
FL03 — Integration Architect - FICO	Focused on cross-module integration: how MM flows into FI, how SD impacts CO-PA, what happens downstream when you change a valuation class. This project contains integration maps and cross-reference tables I've built over time.
FL04 — Specification Writer	For producing functional specifications and technical design documents. It knows my company's document templates, formatting standards, and the level of detail expected. I paste requirements in, and get a structured spec out.
FL05 — Business Process Analyst	For end-to-end process analysis: procure-to-pay, order-to-cash, record-to-report. When I need to document a process flow, identify gaps, or prepare a workshop, this project provides the analytical framework.
FL06 — Change Management & Training Specialist	For the human side of SAP projects: training materials, change impact assessments, communication plans, user adoption strategies. Less technical, more organizational.
FL07 — Proposal & Estimation Assistant	For business development: writing client proposals, estimating project effort, scoping engagements, preparing pitch decks. It knows my rate structure, my past proposals, and my standard engagement templates.

What I learned building this structure

A few things became clear only after months of use:

- **Naming conventions matter.** The FL01-FL07 prefix keeps my projects sorted logically when I scan the sidebar. Without it, projects appear in random order and you waste time finding the right one.
- **Shared files save time.** My Identity Card and Style Guide are loaded in every FL project. I wrote them once; they work everywhere.
- **Not every project gets used equally.** FL01 (Technical Expert) and FL04 (Specification Writer) are my most-used projects. FL06 (Change Management) gets opened maybe once a month. That's fine — it's there when I need it.

MY SKILLS: THE TOOLS I BUILT FOR MYSELF

Skills are where the real leverage lives. A project gives the AI context; a skill gives it a capability. Here are the 6 skills I've built and use regularly.

SAP-Specific Skills

SKILL	WHAT IT DOES
SAP Deep Dive Analyzer	My most-used SAP skill. When I need to understand a business process end-to-end — with table-level detail, cross-module integration points, S/4HANA specifics, and real consulting tips — this skill structures the analysis into 6 layers. It goes far deeper than a standard Claude answer because the skill forces a systematic framework: business walkthrough, technical data flow, cross-module integration table, S/4HANA delta, consulting tips, and further exploration prompts.
SAP Root Cause Investigator	For when something is broken. Period-end closing failures, reconciliation gaps, unexpected account postings, valuation errors. The skill follows a structured diagnostic methodology: symptom capture, hypothesis generation, systematic verification (with specific tcodes and tables to check), and resolution with prevention recommendations. It also searches the web for relevant SAP Notes.

Productivity Skills

SKILL	WHAT IT DOES
Meeting Actions Tracker	My post-meeting workhorse. I paste raw workshop or call notes, and the skill transforms them into structured minutes: key takeaways, a decisions table (with "Decided by" column), an actions table (with owners and deadlines), and an open points parking lot. It matches the language of the input. I trigger it with /meeting-notes.
Prompt Technique Advisor	Instead of knowing by heart all prompting techniques, I use this skill as my trump card. Every time I start a conversation willing to go deep on a subject I'm not familiar with, or to unlock a conversation where I feel stuck, I call this skill. It recommends the best prompting technique for the situation, proposes 3 ranked options with taglines, and then applies

	the chosen technique to rewrite my question. It's a meta-skill — it makes all my other prompts better.
Thinker Prompt Advisor	This one helps me get smarter — in the purest sense of the word. Anytime I need to see a topic from a different angle, with a deeper perspective from an expert, I call this skill. It recommends intellectual frameworks from great thinkers (Socrates, Hegel, Popper, Munger, Feynman, Meadows, Klein) and applies the chosen framework to my current question. When I'm stuck on a strategic decision, this is what gets me unstuck.
Project Instructions Refiner	This one is the least-used skill on my list, but maybe the most useful one. After a long conversation where I had a hard time getting what I wanted and then finally got it right, I call this skill. It learns from the conversation: what went wrong, what finally worked, and why. Then it produces a set of refinements for my project-level instructions so the same friction doesn't happen again. You see the improvement loop? Thanks to this skill, the more I work with AI today, the faster the results tomorrow.

What I learned building skills

- **Start with your biggest time sink.** The Meeting Actions Tracker was my first productivity skill, and it's still the one with the highest ROI. I was spending 20-30 minutes after every workshop formatting notes. Now it takes 30 seconds.
- **Skills compound with context files.** A skill in isolation is useful. A skill running inside a project with loaded context files is powerful. The SAP Deep Dive Analyzer produces dramatically better output when it has my Industry Glossary and client-specific terminology available.
- **Meta-skills are underrated.** The Prompt Technique Advisor and Project Instructions Refiner don't produce deliverables directly. They improve everything else. Building meta-skills was one of the best decisions I made.
- **Iterate relentlessly.** None of my skills were right on the first try. The SAP Root Cause Investigator went through four revisions before the diagnostic methodology matched how I actually troubleshoot problems. Build, test with a real case, fix, repeat.

> **The bigger picture**

> Looking back, the single most impactful thing I did wasn't learning a specific prompting technique or building a specific skill. It was the decision to treat AI as infrastructure rather than a toy. The moment I stopped asking "what can AI do?" and started asking "how do I build a system that makes me better at my job every day?" — that's when everything changed. I hope this guide helps you make that same shift.