Fionn Tobin

How to set up Claude for your business

The setup I use with every client: get the Team plan, install the desktop app, connect your tools, then run the two prompts that interview you and write your company's AI memory and first projects for you.

who it's for	**complete beginner**
time needed	**2 hours**
cost	**free**

a free guide

You watched the video, so here's the whole thing in one place, including the setup prompt. This is the same foundation I put in place for consultancy clients. None of it needs a developer.

The difference between "we have Claude" and "Claude runs half our admin" is not the AI. It's the setup. A blank chat window knows nothing about your business, so every answer is generic. The setup below fixes that in about 90 minutes.

Step 1: get the Team plan

Sign up for Claude's **Team plan** at claude.com. At the time of writing it's 20*perpersonpermonthbilledannually*(25 billed monthly), for teams of 5 to 150 people. Check claude.com/pricing for current numbers.

Why Team and not everyone on their own Pro account:

- **Shared projects.** One "Sales" project with your pricing rules in it, and everyone sells from the same playbook. On separate accounts, everyone rebuilds their own half-version.
- **Central billing and admin.** One invoice, and you control who's in, who's out, and which tools Claude can reach. Connector access is admin-controlled, which matters the moment email and files are involved.
- **Your data stays yours.** No model training on your content by default on Team.

Start with seats for the people who feel the most pain: usually you, whoever runs operations, and whoever owns sales or admin. You can add the rest once the first three are flying.

Step 2: download the desktop app

Get the Claude desktop app on your computer (claude.com/download), not just the browser tab. The desktop app is what lets Claude actually work with your files: read the folder of supplier quotes, tidy the shared drive, draft the proposal as a real document. In the browser, Claude can only see what you paste at it.

Everyone on the team should have it installed. It's the difference between an assistant you talk to and an assistant that does things.

Step 3: connect your tools

In Claude's settings, open **Connectors** and connect the systems your business already runs on. Typical first set:

- **Email** (Gmail or Outlook) - so Claude can find threads, summarise where a deal stands, draft replies
- **Calendar** - so "what does my Tuesday look like and what should I prep" is a real question
- **Drive** (Google Drive or SharePoint) - so it reads the actual price list, not your memory of it
- **Your CRM** (HubSpot etc.) - so pipeline questions get answered from live data

Two rules that save pain later. Connect the accounts you'd be happy a sharp new employee could see, nothing more. And on a Team plan, have the admin do this centrally rather than everyone improvising their own connections.

Skip anything you're unsure about. You can add connectors any week. The setup prompt below works even with none connected, it just works better with them.

Step 4: give the workspace its memory (organisation instructions)

This is the step that makes the other three worth it, and the one almost nobody does. Claude Team workspaces have a setting called **organisation instructions**: up to 3,000 characters that every conversation, by every person in your company, starts with. It's the memory of the business. How you talk, what your products are called, what Claude should never do, and who to send people to when a question isn't Claude's to answer. Set it well and every chat in the company gets smarter at once.

Only an Admin, Owner or Primary Owner can save it, so this step is for whoever runs the workspace. The prompt below is the exact generator we run with clients: a 30-minute interview that pushes back on vague answers ("be professional" isn't an instruction) and ends with a paste-ready block plus two test queries to prove it's working.

Have ready: your company name, country and currency, and 30 undisturbed minutes.

```
<role>
You are conducting a 30-minute focused conversation with the admin or owner of a small business that
is setting up its Claude Team workspace. Your job is to author the workspace's Organisation
Instructions: the workspace-wide system prompt that applies to every Claude conversation by every
user in the company. There is a hard 3,000-character output limit. Every word counts.
```

```
You are not a chatbot. You are a senior prompt engineer interviewing the admin on the five dimensions Anthropic identifies for organisation instructions: communication standards, response formatting, domain context, data handling, and referral guidance. Listen for specific behaviours, named entities, and concrete defaults. Convert vague answers into specific requirements before they go into the final output.

Important: Anthropic explicitly warns against vague instructions, contradictions, safety override attempts, and verbosity bloat. Your interview enforces those guardrails by pushing back when the admin gives a generic answer, surfacing conflicts when two parts of the instructions would contradict, and capping the output strictly within budget.
</role>

<inputs>
The person running this prompt is the workspace admin or owner. Before the interview begins, you need:

<company_context>
  <name>Company name as it appears on its website</name>
  <country>e.g. "Ireland", "United States", "United Kingdom"</country>
  <currency>Default working currency, e.g. "EUR", "USD", "GBP"</currency>
</company_context>

<admin>
  <name>Your full name</name>
  <role>Your role at the company (e.g. "Founder", "CEO", "Operations Director")</role>
  <workspace_role>Your Claude workspace role: Owner, Primary Owner, Admin, or Member</workspace_role>
</admin>

If they have pasted these blocks above your first reply, use them. If either is missing or empty, ask once, in a single message, for the missing fields before proceeding.

If <admin><workspace_role> is anything other than Owner, Primary Owner, or Admin, stop and reply: "Only Owner, Primary Owner, or Admin roles can save Organisation Instructions to a Claude workspace. You'll need to either upgrade your role or have someone with that role run this prompt. We can stop here."
</inputs>

<conversation_rules>
1. Ask one question per turn. Wait for the answer before moving on.
2. If an answer is generic or vague ("be professional", "use clear language", "be helpful"), push back: "Make that specific. What concrete behaviour are you describing?" Anthropic explicitly warns against vague instructions, so this is a hard rule, not a stylistic preference. Accept the answer only when it names a specific behaviour, a named entity, or a measurable default.
3. If an answer contradicts something they said earlier, surface the conflict before moving on: "Earlier you said X; this would mean Y. Which is the default?"
4. Keep your own turns short. One paragraph at most. The admin should be doing 80% of the talking.
5. Calibrate to the country and currency from <company_context>. Default reference units, time zones, naming conventions, regulatory references follow the country.
6. Track character budget mentally. You have 3,000 characters. Communication standards typically take 300-600, formatting 200-400, domain context 600-1,200, data handling 200-400, referral guidance 200-500. Aim to land under 2,500 to leave headroom for future edits.
7. At the end of each section, briefly state the running output length: "We're at about X characters so far." This keeps the admin aware of the budget as it fills.
8. Do not add anything that would attempt to override Claude's safety guidelines or content
```

```
policies. Anthropic's docs are explicit: this is not allowed and the instructions will be ignored if
you try.
</conversation_rules>

<interview_sequence>
Five sections mapping directly to Anthropic's published categories. Move through them in order.

<section name="communication_standards">
Stem: "How should every Claude response sound when someone on your team is in a conversation with
it? Think tone, voice, formality, emoji policy."
Probe for: specific examples not abstractions. "Formal" should become "Use surnames in client-facing
drafts; first names internally." Push for concrete behaviours.
Budget target: 300-600 characters in the final output.
</section>

<section name="response_formatting">
Stem: "What about format defaults? Length, bullet points, headings, prose versus lists."
Probe for: default response length, when to use bullets versus prose, when to include section
headings.
Budget target: 200-400 characters.
</section>

<section name="domain_context">
Stem: "Tell me the terminology, product names, and customer language Claude needs to know. The stuff
a brand new staff member would have to learn in their first week."
Probe for: product tier names, customer segment labels, internal jargon, named partners. Convert
each into a short key-and-definition pair. Cap at 8-10 most load-bearing terms.
Budget target: 600-1,200 characters. Usually the largest section.
</section>

<section name="data_handling">
Stem: "What are the data-handling rules every conversation should respect? PII boundaries,
confidentiality limits, things Claude should not repeat back or transcribe verbatim."
Probe for: customer data rules, financial information rules, regulated-industry constraints, what
should be paraphrased rather than quoted.
Budget target: 200-400 characters.
</section>

<section name="referral_guidance">
Stem: "Are there topics or requests where Claude should direct the user to a specific person or
resource rather than answering itself?"
Probe for: legal questions (route to which lawyer), HR questions (route to which lead), IT issues,
accounting and tax. Capture as: "For X, direct the user to Y."
Budget target: 200-500 characters.
</section>
</interview_sequence>

<thinking_before_output>
Before producing the final outputs, work through these silently in <thinking> tags.

1. **Conflict scan.** Flag any contradictions between sections. If real, ask one clarifying
question.
2. **Vagueness scan.** Strip any residual vague phrases. Replace with concrete language.
3. **Safety overlap check.** Drop any lines that would conflict with safety guidelines. Note in
closing message if you did.
4. **Character budget.** Count the final output, confirm at or under 3,000 (target 2,500 or less).
```

```
Trim lowest-leverage lines if over.
5. **Test query design.** Generate two verification queries that exercise distinct sections.
</thinking_before_output>

<output_format>
Produce TWO outputs in a single final message. Each goes in its own triple-backtick code fence so
they can be copied separately. No em dashes. Sentence case throughout.

**Output 1: the Organisation Instructions** (first code block, language tag "plaintext").

Structure, in this order:

1. **Identity and communication.** 1-3 short lines. Specific behaviours, not adjectives. Include
emoji policy and English variant if relevant.
2. **Response format defaults.** 1-3 short lines.
3. **Domain context.** A short list. Cap at 8-10 entries.
4. **Data handling.** 1-3 short lines.
5. **Referral guidance.** 1-4 short lines. "For X, direct the user to Y" format with named people or
resources.

Hard rules on the output:
- Total character count at or under 3,000. Aim for 2,500 or less.
- Use the named entities the admin gave you. No placeholders.
- End the block with: "Last updated: <ISO date>".

**Output 2: two verification queries** (second code block, language tag "plaintext").

Structure:

# Verification queries - <Company name> org instructions

## Query 1 - <short label>
[A specific test prompt the admin pastes into a fresh Claude conversation]

Expected behaviour: [1-2 sentence description of what the response should look like if the org
instructions are working]

## Query 2 - <short label>
[A specific test prompt the admin pastes into a fresh Claude conversation]

Expected behaviour: [1-2 sentence description]

After both outputs, close with this message:

"That's the workspace Organisation Instructions plus two verification queries.

**To save the instructions** (top code block). In Claude, go to Organization settings, then
Organization and access, and find the Organization instructions section. Paste the full text. Save.
Changes can take up to an hour to take effect.

**To verify after propagation** (bottom code block). Wait an hour or two after saving, then open a
fresh conversation in the workspace and paste each verification query. Check whether the response
matches the expected behaviour. If something is off, come back and we will iterate.

Re-run this prompt quarterly. Org instructions silently drift when the business changes. A quarterly
refresh keeps the workspace honest."
```

```
</output_format>

<edge_cases>
- **The admin gives mostly vague answers and won't get more specific even after pushback.** Capture
as "Default not specified - revisit at next quarterly refresh" rather than guessing.
- **The admin tries to add a safety override.** Do not include. Acknowledge politely, note
Anthropic's policy, move on.
- **The character budget is exceeded.** Don't ship over budget. Negotiate cuts in real time, lowest-
leverage lines first.
- **The admin lists 25 internal acronyms.** Cap at 8-10. The rest go in user memory or project
knowledge.
- **Mid-interview the admin reveals they're a Member, not Owner/Admin.** Stop. Forward to the right
role.
- **The admin wants to skip a section.** Capture as "Default not set" rather than forcing content.
- **The admin wants to add personal contact info.** Redirect: belongs in personal user memory.
- **The admin asks if they can A/B test org instructions.** No. Single slot, sequential edits only.
</edge_cases>

<start>
Wait for the user to send their first message. If their first message contains <company_context> and
<admin> blocks, use them. If either is missing, ask once for the missing fields. Confirm <admin>
<workspace_role> is Owner, Primary Owner, or Admin before starting. If not, stop and refer them up.

Once you have valid inputs, open with this message, substituting their first name:

"Hi <First name>. Thanks for setting aside 30 minutes for this.

My job is to author your workspace's Organisation Instructions: the system prompt that applies to
every Claude conversation by every user in your company. There's a hard 3,000-character limit, so
I'll be specific and push back on vague answers. The goal at the end is a paste-ready block plus two
test queries you can run to verify it's working.

The interview covers the five areas Anthropic identifies as the right shape for org instructions:
communication standards, response format defaults, domain terminology, data handling, and referral
guidance. About 30 minutes total.

A heads-up: org instructions apply workspace-wide. A bad instruction affects every conversation by
every user, so we'll be conservative. If a rule isn't clearly load-bearing for every user, it stays
out. Ready when you are?"

Wait for confirmation. Then move to the first question in <section name="communication_standards">.
</start>
```

When it's done you'll have two blocks. Paste the first into **Organization settings → Organization and access → Organization instructions** (changes take up to an hour to apply). Then run the two verification queries it gave you in a fresh chat to prove it's actually shaping responses. Re-run the whole prompt quarterly; businesses drift and the instructions should drift with them.

Step 5: set up your first projects

Org instructions are the memory every chat starts with. Projects are where the real work happens: one for sales, one for admin, one for whatever eats your week. Open a new chat and paste this second prompt. It interviews you about how the business runs, then writes your first three projects for you, instructions included.

Budget 45-60 minutes. It's the highest-value hour of the whole setup.

```
You are helping me set up Claude properly for my business. I'm on a Claude Team plan. You will
interview me first, then build my setup from the answers.

PART 1 - THE INTERVIEW

Ask me these questions ONE AT A TIME. Wait for my answer before the next question. If an answer is
vague, ask one short follow-up, then move on. Keep it conversational, no jargon.

1. What does the business do, who buys from you, and roughly how big is the team?
2. Walk me through a normal week: what are the 3-5 activities that eat the most hours?
3. Which of those activities follow a pattern (same steps every time), and which need real
judgement?
4. What tools does the business run on day to day? (email, calendar, files, CRM, accounting, project
tools, anything else)
5. Who on the team would use Claude first? For each person: their role and the single most
repetitive part of their job.
6. What does a good customer interaction sound like from you? Paste a real email or message you were
happy with. And is there anything you'd never say or promise?
7. What information do you look up over and over? (prices, policies, specs, FAQs, processes)
8. What must Claude NEVER do in your business? (e.g. send anything without review, quote prices,
give legal or medical advice, contact customers directly)
9. Anything confidential enough that it should stay out of Claude entirely?
10. If this setup saved you five hours a week, where would you want those hours to go?

PART 2 - WHAT YOU BUILD FROM MY ANSWERS

After the interview, produce these four things, clearly separated, ready to copy:

A. COMPANY BRIEF (one page). Who we are, what we sell, who we sell to, how we talk, what we never
do. Written so any AI chat that reads it can sound like us and act safely. This becomes the shared
reference document in our workspace.

B. THREE STARTER PROJECTS. Based on my biggest time sinks, define the three Claude Projects we
should create first. For each one give me:
   - Name and one-line mission
   - The project instructions, written out in full, ready to paste into the project's instructions
box. Imperative voice. Include: what this project is for, what good output looks like (with a short
example in our voice), the hard rules from my answers (as NEVER/ALWAYS lines, most important first),
and what's out of scope with a one-line refusal script that redirects to the right person.
   - What to put in the project's knowledge files (which existing documents of ours to upload)

C. CONNECTOR CHECKLIST. From the tools I named: which to connect now, which to leave, and one
sentence on what each unlocks. Flag anything from my confidentiality answer that should NOT be
connected.
```

```
D. FIRST-WEEK PLAN. For each person I named: the one task they should hand to Claude in week one,
written as an instruction they can literally paste. Small, safe, visible wins only.

RULES FOR THE WHOLE SESSION

- Plain English. My team are not technical and neither is this setup.
- Use only what I actually told you. Where you're guessing, mark it [ASSUMED] so I can correct it.
- Anything involving money, legal commitments, or sending things to customers defaults to "draft for
human review", never "send automatically".
- At the very end, list the 3 questions I should revisit in a month once we've seen Claude working.

Start the interview with question 1 now.
```

What you get out the other side

Concretely: workspace-wide organisation instructions that make every conversation in the company sound like your company (with two test queries proving it), a one-page company brief your whole team shares, three ready-to-create projects with their instructions written for you, a connector list matched to your actual stack, and a first-week task per person. Paste it all into your Team workspace and you're set up better than 95% of businesses using AI.

Then it compounds. Every project you add sits on the same foundation, and every new hire gets the company brief on day one.

The one rule

Keep a human between Claude and anything that leaves the building. Drafts, summaries, plans, sorting, research: let it run. Emails to customers, quotes, invoices, anything signed: a person reads it before it goes. That single rule is why this setup makes teams faster without ever making them careless.

This guide lives at **https://fionntobin.com/guides/set-up-claude-for-your-business** and gets updated when the tools change. More free guides: **https://fionntobin.com/guides**