



Setting Up Claude for Your Business

Which plan, which model, what it really costs, and how to roll it out — a no-fluff setup guide for owners who want AI working in the business, not just in a browser tab.

- › Every plan compared — Free to Enterprise
- › Opus vs. Sonnet vs. Haiku, in plain English
- › Real usage costs & the levers that cut them
- › A step-by-step rollout you can run this week

The Claude Setup Playbook

2026 EDITION • ~12 MINUTE READ

FREE GUIDE

Two ways to put Claude to work

Before you pick a plan or a model, get this one decision right. Almost every business uses one or both of these paths — and they're billed completely differently.

PATH A — SUBSCRIPTION

The Claude apps

A flat monthly fee for you and your team to use Claude through the chat app, the **Cowork** desktop agent, and **Claude Code**. No engineering required. This is how most owners and teams start.

Best for: day-to-day work — writing, research, analysis, planning, building docs and decks, running agents on your files.

PATH B — API / PLATFORM

Build it into your stack

Pay only for what you use, billed per **token**, to wire Claude into your own products, automations, and workflows (e.g. a support bot, lead enrichment, document processing).

Best for: automation and software — anything that needs Claude running behind the scenes at scale.

THE ONE THING OWNERS GET WRONG

A subscription does **not** include API access, and the API does not include the apps. They're separate products on separate bills. Most businesses run a subscription for the team *and* a small API budget for automations — that's normal and expected.

How to read the rest of this guide

If you mainly want to...	Go to	Why
Get your team using Claude	Section 2 — Plans	Pick the right subscription tier and seat count.
Know which "brain" to choose	Section 3 — Models	Opus, Sonnet, and Haiku each have a job.
Understand & control spend	Section 4 — Costs	How tokens work and how to cut the bill.
Actually roll it out	Section 5 — Setup	A step-by-step launch checklist.

Subscription tiers, decoded

The plan names don't tell you what's different — so here it is in plain terms. Pro to Max buys *usage headroom*, not new features; Team and Enterprise add admin, security, and collaboration.

Plan	Price	Who it's for	What you get
Free	\$0	Trying Claude; light personal use	A capable model with a daily message cap. No Cowork or Code; lower priority when busy.
Pro	\$20/mo \$17/mo billed yearly	Solo owners & professionals — the default starting point	~5× Free usage, all models, Projects, plus Claude Code and the Cowork desktop agent.
Max 5×	\$100/mo	Power users who hit Pro's limits	5× the usage of Pro and higher priority. Same features as Pro.
Max 20×	\$200/mo	All-day heavy users & agent workflows	20× Pro usage and top priority. Same features as Pro.
Team	~\$30/seat/mo ~\$25 yearly · 5-seat min	Small businesses & departments	Everything in Pro + central billing, admin controls, SSO, shared Projects, and no training on your team's conversations by default. Premium seats (~\$150/mo) add ~5× usage.
Enterprise	Custom	Larger orgs & regulated industries	Advanced security & compliance, data-residency options, more usage, and dedicated support.

WHY OWNERS LIKE SUBSCRIPTIONS

- + Predictable flat cost — no surprise bills
- + Zero setup; your team can start today
- + Team tier keeps your data out of model training by default

WATCH-OUTS

- Usage limits reset on a window — heavy days can hit caps
- Team requires a 5-seat minimum commitment
- Doesn't cover anything you want to build with the API

RULE OF THUMB

Start on **Pro**. Hitting limits mid-afternoon? Move to **Max 5×**. Once 3+ people use it, switch to **Team** — similar per-seat cost to Pro, but you get governance and privacy controls.

Opus, Sonnet & Haiku — which brain?

Claude comes in three current-generation models — think of them as a team: a heavyweight strategist, a fast all-rounder, and a high-speed specialist. On the apps you switch freely; on the API you choose per request.

Model	Think of it as...	Reach for it when...	API price /M tokens
Opus 4.8 Flagship	Your most capable strategist. Best at hard reasoning, long multi-step work, and complex coding.	The problem is genuinely hard or high-stakes and quality matters more than speed.	\$5 in \$25 out
Sonnet 4.6 Balanced	The dependable all-rounder. Strong intelligence at a far lower cost — the everyday default.	Most real work: writing, analysis, summaries, tool use, day-to-day coding.	\$3 in \$15 out
Haiku 4.5 Fast	The speed specialist. Quick and cheap, great at high-volume, well-defined tasks.	You're running lots of simple jobs: classifying, extracting, routing, tagging.	\$1 in \$5 out

OPUS 4.8 — PROS & CONS

PROS

- + Top reasoning & coding
- + Handles long, complex tasks

CONS

- Highest per-token cost
- Overkill for simple work

SONNET 4.6 — PROS & CONS

PROS

- + Best value for quality
- + Fast enough for most work

CONS

- Not quite Opus on the hardest tasks

HAIKU 4.5 — PROS & CONS

PROS

- + Cheapest & fastest current model — ideal for volume
- + Great for automations that fire constantly

CONS

- Less depth on nuanced reasoning
- Best kept to clearly-scoped tasks

THE MONEY MOVE

All three share the same large context window and 5x-output-to-input pricing. Don't run everything on Opus — **default to Sonnet, drop to Haiku for high-volume simple jobs, and escalate to Opus only for the hard ones.** This alone often halves API spend.

What it actually costs to run

Subscriptions are simple — a flat monthly fee. The API is where owners get nervous, so let's demystify it. You're billed per **token**, separately for input and output.

WHAT'S A TOKEN?

Roughly **4 characters**, or about $\frac{3}{4}$ of a word. Everything you send (your prompt, instructions, history, attached docs) is **input**; everything Claude writes back is **output**. Output costs 5× input — so concise answers are cheaper.

HOW YOU'RE BILLED

The API is pure pay-as-you-go on prepaid credits in the Console. Send nothing this month, owe nothing. There's no monthly minimum and no permanent free production tier — you fund credits and draw them down.

Real-world monthly examples

A moderately busy automation using ~2M input + 500K output tokens per month, before any discounts:

Model	Approx. monthly cost	Good fit for this budget
Haiku 4.5	~\$4.50	High-volume tagging, routing, extraction
Sonnet 4.6	~\$13.50	A capable support or ops assistant
Opus 4.8	~\$22.50	Complex reasoning where quality pays off

Costs scale linearly with volume — 10× the usage is roughly 10× the cost. Estimate your own with: $(input\ tokens \div 1M \times input\ rate) + (output\ tokens \div 1M \times output\ rate)$.

Four levers that cut the API bill

1 • Route by difficulty

Send simple tasks to Haiku, everyday work to Sonnet, hard problems to Opus. The single biggest saver.

2 • Prompt caching

Reuse the same system prompt or documents? Caching can cut that repeated input cost by up to ~90%.

3 • Batch processing

For non-urgent jobs (overnight runs, bulk processing), the Batch API gives a flat 50% discount.

4 • Trim the context

Don't paste in everything. Less unnecessary input = fewer tokens = a smaller bill on every call.

FOR MOST OWNERS

You'll likely run a **subscription** for the team's hands-on work, plus a small **API budget** (often tens of dollars, not thousands) for a few automations. Set spend alerts in the Console.

Set up Claude for your team

A clean rollout you can run this week. This is the subscription path — the right starting point for almost every business.

1

Pick your tier and claim the account

Solo? Start on Pro. A team of 3+? Go straight to Team for admin and privacy controls. Sign up at claude.com and choose monthly or (cheaper) annual billing.

2

Lock down privacy & data settings first

On Team/Enterprise, your conversations aren't used for model training by default — confirm it in admin settings. Decide what data is and isn't allowed in prompts, and write a one-page internal AI policy.

3

Invite your team & set roles

Add seats, assign an admin or two, and group people sensibly. Keep an eye on the 5-seat minimum on Team.

4

Build Projects for recurring work

Projects let you save context, instructions, and reference files in one place so Claude "knows" a workflow. Create one per function — Marketing, Sales, Ops, Support — and load each with your brand voice, SOPs, and templates.

5

Connect your tools

Use connectors to plug Claude into the apps you already run (docs, calendars, knowledge bases, and more). Install the Cowork desktop app for agentic work on your own files, and Claude Code if anyone touches development.

6

Seed 3 high-value use cases & train the team

Don't boil the ocean. Pick three painful, repetitive tasks, build a shared prompt or Project for each, and run a 30-minute team walkthrough. Wins compound from there.

PRO TIP

The teams that get value fastest treat prompts like SOPs: write them once, save them in a Project, and reuse them. A good saved prompt is a reusable asset — not a one-off question.

If you're automating with the API

For owners who want Claude running behind the scenes. You don't have to do this yourself — but here's the shape of it so you can brief whoever does.

1

Create a Console account & fund credits

Set up the developer Console, add prepaid credits, and turn on spend limits and alerts before anything goes live.

2

Choose your default model

Start on Sonnet for quality-per-dollar. Drop to Haiku for high-volume steps; reserve Opus for the genuinely hard ones.

3

Build one workflow, measure, then scale

Ship a single automation (e.g. lead enrichment or inbox triage), watch the token cost per run, then turn on caching/batching and expand.

Where businesses get the most value

Marketing

Ad copy, content repurposing, campaign briefs, landing-page drafts, audience research.

Sales

Personalized outreach, call-note summaries, proposal drafts, CRM cleanup, lead qualification.

Operations

SOP writing, process docs, meeting summaries, data cleanup, and internal Q&A on your own files.

Customer support

Draft replies, build help docs, triage tickets, and power an on-brand support assistant.

Finance & admin

Spreadsheet analysis, invoice/contract review, report drafting, and turning data into plain-English summaries.

Leadership

Strategy sparring, scenario planning, board-deck drafts, and digesting long documents fast.

The 60-second cheat sheet

Stuck on a choice? Match your situation to the answer.

Just exploring on my own?

Start **Free**, upgrade to **Pro (\$20/mo)** the moment it's useful.

Hitting limits as a solo user?

Move to **Max 5× (\$100/mo)**; go **20× (\$200/mo)** only if you're in it all day.

Rolling out to a team of 3+?

Team — for admin, SSO, shared Projects, and no-training-by-default privacy.

Regulated, or 50+ people?

Talk to **Enterprise** for compliance, data residency, and dedicated support.

Which model day-to-day?

Sonnet 4.6 by default. **Haiku** for volume, **Opus** for the hard stuff.

Want Claude inside my software?

Use the **API** — pay per token, start with a small credit budget and alerts.

Worried about cost?

Route by difficulty, cache repeated context, batch non-urgent jobs, trim input.

THE 3-LINE STARTING PLAN FOR MOST OWNERS

1. Pro or Team subscription for the people. 2. Sonnet as the everyday model. 3. A small API budget for one or two automations once the team is comfortable.

YOUR NEXT STEP

You've got the map. Now **make the first move.**

You don't need a perfect plan to start — you need one workflow working. Here's your launch checklist:

- ✓ Pick your tier — Pro to test, Team to roll out
- ✓ Set Sonnet as your default model
- ✓ Build one Project for your most repetitive task
- ✓ Run a 30-minute team walkthrough
- ✓ Add a small API budget only when you're ready to automate



YOUR GUIDE

Jason Mance

Omega Tecks · AI Strategy & Implementation



WANT THIS TAILORED TO YOUR BUSINESS?

Let's map your AI setup together.

[Book My Free Session →](#)