



FOUNDERS FIELD KITS · ONBOARD YOUR CLAUDE · KIT 03

Claude, Locked Down



Every privacy setting in plain English. What Claude can see, what it can't, what "training" actually means, and how to keep client information exactly where it belongs. Fifteen minutes, fully sorted.

FOUNDERS *keeper*

OPERATIONS · AI · AUTOMATION

The chapter of AI nobody explains (so everyone quietly worries)

Every founder I've onboarded to Claude has asked me some version of the same question, usually in a lowered voice: *"But... where does everything I type actually go?"*

Fair question. Great question, actually. You handle client information, financials, maybe health or legal details. "Just trust the AI company" is not a privacy policy, and the actual privacy policies read like they were written to be skimmed and surrendered to.

So here's the deal. This kit walks through **every setting that matters, in plain English**: what it means, what I recommend, and why. By the last page you'll have a locked-down account and, more valuable, **a simple rule for what you paste and what you never paste**, so you can use Claude confidently for the rest of your business life.

- ☐ Understand who can see your chats (the honest version)
- ☐ Set the training toggle deliberately instead of by default
- ☐ Adopt the traffic-light rule for client data
- ☐ Know exactly what a connector can and can't touch
- ☐ Know how to export or delete everything, any time
- ☐ Tick every box on the printable Lock-It-Down checklist

Keeper's note

Settings screens move around and policies get updated. This kit is accurate as I write it (July 2026), and the *principles* don't date. If a toggle in here has moved, it exists under a nearby name. When in doubt, Anthropic's privacy centre at privacy.claude.com has the current word.

Chapter one

Who can actually see your chats

Let's kill the two extreme beliefs first, because both are wrong:

- **"It's totally private, like my diary."** Not quite. Your chats sit on Anthropic's servers, like your emails sit on Google's.
- **"Someone's reading everything I type."** Also no. Nobody is sitting there reading your invoice reminders.

The honest picture:

- **You** can see your chats, in your account, until you delete them.
- **Anthropic's systems** process them (that's how Claude answers) and automated safety systems scan for misuse. Human review is the exception, tied to things like safety flags, not routine browsing.
- **Other users can't see them. Ever.** Claude will not leak your pricing strategy to the next person who asks. Your chats are not floating in some shared pool.
- **Anthropic doesn't sell your conversations to advertisers.** Their business model is subscriptions, not ads. This is one of the reasons I recommend Claude to clients in the first place.

The one decision that actually matters: training

The big question is whether your chats can be used to **train future versions of Claude**, and that one is your choice, via a setting. Which brings us to chapter two.


SHORTCUT SUMMARY

- Chats live on Anthropic's servers, **visible to you, not to other users**
- No humans routinely reading; **no selling to advertisers**
- The decision you control: **the training setting** (next chapter)

Chapter two

The settings tour: what I recommend and why

Open claude.ai → click your avatar → **Settings**. We're heading to the Privacy section. This whole tour takes ten minutes.



The Privacy page: protection info up top, the preferences that matter in the middle, your data controls below.

Setting 1: "Help improve our AI models" (the training toggle)

Preferences

Location metadata

Allow Claude to use coarse location metadata (city/region) to improve product experiences. [Learn more.](#)



Help improve our AI models

Allow the use of your chats and coding sessions to train and improve Anthropic AI models. [Learn more.](#)



The two that matter, up close: location metadata and the training toggle. Both off is my recommended resting state.

What it means in laymen's terms: if this is ON, your conversations can be used as study material for future versions of Claude, and they're kept for years to make that possible. If it's OFF, your chats are only kept for the normal short retention window and aren't used to teach the model anything.

My recommendation for business owners: turn it OFF. Not because anything sinister happens when it's on, but because you handle other people's information. Client details are not yours to donate to science. Off means the question never arises.

Keeper's note

You were asked about this when you signed up, probably on a screen you clicked through at speed like every human alive. Today's the day you set it on purpose.

Setting 2: Chat history and memory

Claude keeps your conversations so you can come back to them, and (on paid plans) can carry memory between chats so it remembers context you've shared. Handy, and worth knowing it's there. If a chat contains something sensitive, delete that chat when you're done (chapter five shows you how) rather than letting it sit around out of habit.

Setting 3: The little extras ("Location metadata" and friends)

Anything opt-in that shares extra data beyond what's needed to run the product, like the location metadata toggle sitting above the training one: my default is off. You can always turn things on later once you know why you'd want to.

Chapter three

The traffic-light rule for client data

Settings protect you from the platform. This chapter protects you from *yourself at 9pm, tired, pasting an entire client email thread without thinking*. One rule, three colours, zero decisions required in the moment:

● GREEN · PASTE FREELY

Your own business content: your copy, your policies, your pricing, your brain-dumps, your drafts. Public information. Anything already on your website. This is 90% of daily use and it's all yours to share.

● AMBER · ANONYMISE FIRST (TAKES 10 SECONDS)

Client situations where the *story* matters but the *identity* doesn't. Swap names for roles before pasting: "Sarah from Lighthouse Physio" becomes "a physio client". Claude gives you the same quality answer about "a client who's 60 days overdue" as it does about a named human. The advice doesn't need the name. So don't give it the name.

● RED · NEVER PASTE

Passwords, banking details, credit card numbers, Medicare/TFN-style government identifiers, health records, legal matters under privilege, anything you've signed an NDA about, and anything a client sent you in explicit confidence. Not because Claude will misuse it, but because it doesn't need it, and information that's never shared can never leak. This rule applies to every tool you use, not just AI.

Keeper's note

Print this page and stick it near your screen. In my experience the traffic-light rule is the single thing that turns "I'm scared to use AI for client work" into "I use AI for client work every day, safely". It's also a great little policy to show clients if they ask how you use AI. Instant trust.

One more thing, and it matters: the traffic-light rule is your floor, not your ceiling. Every business should set its own privacy and security policy, and some industries set the bar higher for you. If you work with health, legal, financial or education information, or under confidentiality agreements, check what your regulator, industry body and client contracts expect around AI use, and disclose your use of it where disclosure is due. When you're not sure, that's a question for professional advice, not a guess.


SHORTCUT SUMMARY

-  **Yours?** Paste away
-  **Client story?** Strip the names first
-  **Credentials, IDs, health, legal, NDA?** Never

Chapter four

Connectors: what "Claude can access your..." really means


At some point Claude will offer to connect to your other tools: calendar, email, files. Genuinely useful ("what's my week look like?"), and this is exactly where you should slow down and read, because you're granting access to real accounts.



The Connectors page: everything with access, in one list, each one disconnectable. (This is a well-used account; yours starts empty.)

The plain-English version of how it works

- **Nothing connects by itself.** Every connection starts with you clicking approve on a permission screen. No approval, no access. Silence means no.
- **Access is scoped.** A calendar connector sees your calendar. It doesn't get the run of your laptop. The permission screen lists exactly what's included, and that list is the contract.
- **Read vs act.** Some connections just *look* (read your calendar); some can *do* (send an email, create an event). Know which one you're approving. "Can do" deserves more caution than "can see".
- **Every connection is reversible.** The Connectors page shows everything with access, and disconnecting is one click. No breakup conversation required.



Inside one connector: every tool it can use, itemised, with your choice per tool. Notice the split between interactive tools (can do) and read-only tools (can see). This screen is the contract.

My rule of thumb

Connect tools when a repeating task justifies it, not because you can. Start with zero connectors. Add your calendar when you're actually asking about your schedule weekly. Add email when you're actually triaging with Claude. Each connection should earn its place. An audit of "what has access to what" every few months is a beautiful habit (there's a line for it in the final checklist).

SHORTCUT SUMMARY

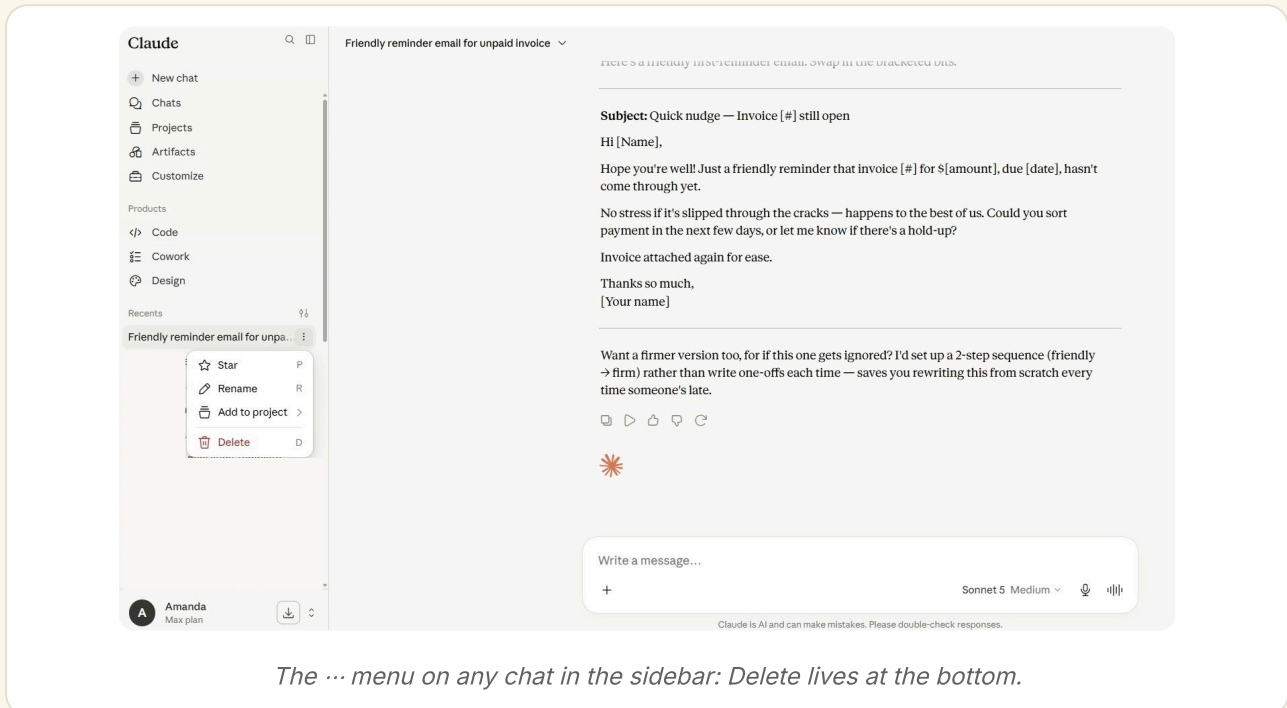
- Connectors only exist if **you approve them**, and the permission screen is the contract
- Know the difference: **can see vs can do**
- Start with none, **add when a real task justifies it**, disconnect any time

Chapter five

Your data, your call: deleting and exporting


Deleting a single chat

Sidebar → hover the conversation → the ... menu → **Delete**. Gone from your view immediately, and removed from Anthropic's systems within their stated deletion window (around 30 days, as with most cloud services). Do this as a habit for any chat that carried amber-level client content.



Exporting everything


Settings → Privacy (or Account) → **Export data**. You'll get an email with a download of your conversations. Worth doing before any big clean-out, and simply good practice: your thinking is business IP, and now there's a copy that belongs to you.



Export: pick the date range, click Export, and a download link lands in your email (it expires after 24 hours, so grab it same-day).

The nuclear option

Deleting your whole account (Settings → Account) removes your data per the privacy policy. You'll almost certainly never need this, but knowing the exit exists changes how comfortable you are inside. Same reason you check where the door is in a new building.



The Account page, simplified: your email, the log-out-everywhere button, and the exit door. (Illustration; your screen's wording may differ slightly.)

The main event

The Lock-It-Down checklist

Fifteen minutes, once, and then you never have to think about this again. Tick as you go:

- ☐ **Training toggle set deliberately.** Settings → Privacy → "Help improve our AI models" → my recommendation: off
- ☐ **Location metadata set deliberately.** Same page, the toggle just above → my recommendation: off
- ☐ **Any extra data-sharing or analytics options reviewed** and set to off unless you have a reason
- ☐ **Traffic-light rule printed** (or screenshotted) and adopted: green paste, amber anonymise, red never
- ☐ **Connectors audited:** Settings → Connectors → everything listed is something you remember approving and still use
- ☐ **One sensitive old chat deleted**, just so you've done it once and know the moves
- ☐ **Data export run once**, so you know what leaving looks like
- ☐ **Phone app checked:** same account, same settings apply, nothing extra to do
- ☐ **Calendar reminder set for a 5-minute re-audit in 6 months** (settings drift, businesses change)

Keeper's note

If a client ever asks "is my information safe with your AI setup?", you can now answer in one breath: *"My account doesn't feed the training data, client details get anonymised before they go anywhere near it, and identifying or financial information never goes in at all."* That sentence is worth more than this kit cost. Use it.

Locked down. Now level up.

Here's the thing about trust: now that you know exactly what Claude can and can't see, you can stop using it at arm's length. The founders who get the most out of AI aren't the ones being reckless with it. They're the ones who sorted the boundaries early and then went all in *inside* those boundaries.

Kit 01 • Claude From Scratch

If you landed here first: the ground-zero setup guide. The screen decoded, the vocabulary sorted, and 10 starter prompts that earn their keep.

Kit 02 • Claude, Meet My Business

The one that changes your week. Fill-in-the-blank templates that teach Claude your business, your customers, and your tone of voice, loaded into a Project so every chat starts already knowing you. Everything you paste into it follows the traffic-light rule you now know by heart.

Keeper's note

You've drawn the boundaries; now put them to work. The **Mental Load Mini Audit** is my deeper dive into where your hours actually go: what's quietly draining you, and exactly which tasks could be automated with Claude and the tools you're already paying for, all inside the rules you set today. You bring the chaos, I bring the map. Grab it at founderskeeper.com.au/services-page/mental-load-mini-audit. Let's talk.



FOUNDERS *keeper*

OPERATIONS · AI · AUTOMATION

© 2026 Founders Keeper · founderskeeper.com.au · Settings and policies described are accurate at July 2026; the current word always lives at privacy.claude.com. This kit is general guidance, not legal advice. For the buyer's use, please don't redistribute (but do tell your business besties where you got it).