




Putting AI to work, safely

A plain-language guide for any team. How to think about AI, which tool fits which work, what skills and connectors are, and a step-by-step path to first results that keeps client data out until you decide otherwise.



How AI fits into a desk at work. AI never talks to a client or changes a record on its own. A person always stands in between.

By Carlos Masson, Forward Deployed Engineer

For owners, managers and every team

01 Ground rules come first

03 Which tool for which work

05 Skills

07 Build, buy, or leave alone

09 A path, one step at a time

02 AI fluency: the 4D framework

04 Getting the most out of Claude

06 MCP connectors

08 Your website and AI answers

10 Free resources

Ground rules come first

- ✓ **The short version.** Use one approved tool, sign in with your company account, keep client data out, and have a person sign off on everything.

Before anyone uses AI for company work, a few decisions need an owner. None of them are technical surprises. They're the same controls you'd want for any new vendor.

	Public and internal	EXAMPLES Office hours, the employee handbook, training material, a draft memo	THE RULE OK in the approved tool
	Confidential	EXAMPLES Budgets, board materials, vendor contracts, internal reports	THE RULE Only in approved, configured tools, under business terms
	Client information	EXAMPLES Names, account or policy numbers, Social Security numbers, health or financial records	THE RULE Stays out until someone accountable approves a specific use

The data traffic light. Before pasting anything into an AI tool, ask which light it is.

- ✓ **One approved tool, on a business plan**
Business plans such as Claude Team or Enterprise come with admin controls and single sign-on, and Anthropic's [commercial terms](#) say customer content isn't used to train its models. Confirm the current terms before you sign.
- ✓ **Sign-in through your company accounts**
Single sign-on and multi-factor authentication, so access ends the day employment does.
- ✓ **A data rule everyone can repeat**
Public and internal information is allowed. Confidential information goes only into approved, configured tools. Client information stays out until someone accountable approves a specific use.
- ✓ **A person signs off**
AI drafts. A named employee reviews and owns anything that leaves the desk.
- ✓ **Permissions set on purpose**
Tools like Claude Cowork reach only the folders and connectors they are given. Decide those limits during setup, not after.



Know where the data goes

By default, Anthropic deletes business API data within 30 days. For work that touches client information, Anthropic offers [Zero Data Retention](#) on its API by agreement, so prompts aren't stored after the answer. It doesn't cover the standard Team or Enterprise chat apps, so decide this during setup.



A one-page AI use policy

Short enough that every employee reads it and can follow it.

HOW TO READ THE EXAMPLES

Every example in this guide uses public or internal information only. Anything involving client data is marked as a later step that needs review.


02 · AI FLUENCY

The 4D framework



The short version. Decide what to hand off, ask clearly, check what comes back, and own the result.

The strongest free foundation available is Anthropic's AI Fluency framework, developed with Professor Joseph Feller (University College Cork) and Professor Rick Dakan (Ringling College). It names four skills, the 4 Ds. We teach this framework rather than inventing our own, so your team learns the same vocabulary as Anthropic's free courses.



The 4 Ds on one task. Delegation happens before you start. Description and Discernment repeat until the draft is right. Diligence covers the whole job.

× **VAGUE**

Write a holiday notice.

Generic text, wrong tone, one language, maybe mentions a sale.

✓ **CLEAR**

Write an office notice for clients in English and Spanish, under 60 words each, saying we're closed on Veterans Day, November 11. Warm tone. No prices or promotions.

Goal

Audience

Format

Limits

Ready to review in one try.

Description in practice. The clear request names the goal, the audience, the format and the limits.

THE CONVERSATION LOOP What you ask for, and how you judge what comes back.

D Description

Say clearly what you want: the goal, the audience, the format, and an example.

AT WORK

Instead of "write a holiday notice," try "Write an office notice in English and Spanish saying we're closed on Veterans Day, November 11. Warm, under 60 words each, no prices or promotions."

Would a new employee understand this request without asking me anything?

D Discernment

Judge what comes back for accuracy, tone, and what's missing.

AT WORK

Every number, date, name and claim is checked against its source before it leaves the desk. AI can sound confident and still be wrong.

What in this would I have to defend, and have I checked it?

THE RESPONSIBILITY LOOP What you hand off, and how you stay accountable for it.

D Delegation

Decide what AI should do and what stays with people.

AT WORK

Good to hand off: first drafts, summaries of long documents, the Spanish version of a notice. Stays with people: decisions about money, people and contracts, and anything that commits the company.

If this comes back wrong, who is affected, and would we catch it?

D Diligence

Use AI responsibly and openly, and own the result.

AT WORK

Only approved tools. Never client data in an unapproved one. Say when AI helped draft something, and keep a person's name on the final version.

Would I be comfortable explaining how this was made to a client?

FREE PRE-WORK

Anthropic's [AI Fluency for Small Businesses](#) course teaches the 4 Ds with business examples in about 14 short, self-paced modules and ends with a certificate. If your team takes it before a workshop, the live session can go straight to real work.

Which tool for which work

✓ **The short version.** Chat for quick questions. Cowork for multi-step work in your files. Code for building tools.

Anthropic offers three ways to work with Claude. Most staff need the first two.



Chat

1 • steps it takes by itself

You ask, it answers. One step.
You copy what you need.

EVERYONE



Cowork

4 • steps it takes by itself

You give a goal. Opens your files, compares, writes, checks. You watch each step.

MOST OFFICE STAFF



Code

6 • steps it takes by itself

You describe a tool. Builds software that then runs every day.

IT, WITH SUPPORT

Does less on its own Does more on its own

Three tools, one difference: how many steps Claude takes before a person looks. More steps means more setup and more oversight.

	WHAT IT DOES	BEST FOR	EXAMPLE
Chat	A conversation. You ask, it answers. Everyone	Quick questions, first drafts, rewording, translation	Rewriting a client email in plain Spanish
Cowork	You give it a goal. It works through your files and tools step by step, and you see each step. Paid plans • reaches only what you allow	Multi-step work across documents: reviewing, reconciling, building a report	Turning a folder of meeting notes into a monthly summary
Code	An agent that builds software and automations. IT, and subject experts with support	Internal tools, scripts, website changes	A small internal tracker your team actually uses

Coding background matters less than people expect. Anthropic [studied about 400,000 Claude Code sessions](#) and found that in each of the ten largest occupations in the data, success rates landed within seven points of software engineers. What predicted success was knowing the subject well.

04 • POWER FEATURES

Getting the most out of Claude

- ✓ **The short version.** Most people only use the chat box. These features are already included in business plans.

Most teams use Claude like a search box: one question, one answer. The features below change how much gets done. Most need no technical skill, though an administrator may have to switch some on first.

FOR EVERYONE • IN THE CLAUDE APP



Projects

A shared workspace with your documents and standing instructions, so you don't re-explain every time.

TRY IT

An "Operations" project with the procedures manual and the style guide.



Research

Searches many sources for a few minutes and returns an answer with citations.

TRY IT

"What are small businesses in Miami doing with AI for customer service? Cite sources."



Connectors

Lets Claude search Outlook, SharePoint, Teams or Google Drive with your own permissions. Read-only by default.

TRY IT

Find the latest version of the employee handbook in SharePoint.

Admin turns on • email after review



Files and Office

Claude creates Excel, Word, PowerPoint and PDF files, and works inside Excel and PowerPoint.

TRY IT

Turn these meeting notes into a 5-slide summary deck.

Check every calculation



Cowork and scheduled tasks

Give it a goal. It writes a plan, works through your files, and shows each step. It can repeat on a schedule.

TRY IT

Every Monday at 8, summarize last week's internal reports into one page.

Paid plans




Memory

Remembers your preferences between chats. You can see and edit what it keeps.

TRY IT

It learns that you want Spanish drafts to use "usted".


Admin turns on




In the workshop, each person tries Projects, Research and the plan step on a real task from their own desk.

Plan mode, in one picture

× WITHOUT A PLAN




✓ WITH A PLAN




Plan mode in Claude Code, and the plan step in Cowork, put your review before the work instead of after it. Fixing a plan takes a minute. Fixing finished work takes an afternoon.

Many agents at once



Subagents and ultracode in Claude Code. Each helper reads its own share, so a big job finishes faster. More agents use more of your allowance, so save this for large jobs.

Make your allowance last



Usage depends on how long the conversation is, the size of attachments, Research, the model and the effort setting. These habits come from Anthropic's usage guidance; the last one is ours.

Free open-source tools: check before you install

Some free tools plug into Claude Code and claim to cut token costs, for example by giving Claude a map of a codebase. Most of those claims come from the tools' own makers, and these tools help developers, not chat users. Run any of them through this check first.



License

Read the license file itself, not the badge. Send GPL or AGPL licenses to legal.



Security record

Published advisories, a security policy, and an OpenSSF Scorecard check.



Agent risks

Prompt injection and tool poisoning, from the OWASP lists for AI and MCP.



Still maintained

A recent release, and more than one person keeping it alive.



Where data goes

Telemetry, cloud sync or calls to another AI service. Can each be switched off?



Pilot and measure

Cost per accepted result on your own work, not the maker's benchmark.

The same review you apply to any vendor applies to free software too.

05 · SKILLS

Skills: teaching AI how your company does things



The short version. A skill is a written playbook, so everyone gets the same good result.

A skill is a written playbook Claude follows whenever a task matches it. It's a folder with a SKILL.md file of instructions, plus any templates, checklists or reference documents the task needs. Claude reads the full skill only when it's relevant, so a team can keep many. Skills follow an [open standard](#), so they aren't locked to one product.

The person who knows the job best writes it down once, and everyone gets that person's standard every time.

WITHOUT A SKILL




Same request from
three people



× 120 words
× English only



× Mentions a promotion
× Wrong tone




× No closing line
× 45 words

WITH A SKILL



Same request from
three people



✓ EN + ES
✓ under 80
words
✓ approved
closing



✓ EN + ES
✓ under 80
words
✓ approved
closing



✓ EN + ES
✓ under 80
words
✓ approved
closing

Why skills matter. Without one, every person gets a different result. With one, everyone gets your company's standard.

office-notice/SKILL.md

example

name: office-notice

description: Writes client-facing notices in English and Spanish in our voice. Use for closures, hours changes and events.

1. Write both languages. Spanish uses "usted" and « » quotes.
2. Keep each version under 80 words.
3. Never mention prices or promotions.
4. Close with an approved line from closing-lines.md.
5. Label the result "Draft · needs review".

Brand voice check

Reviews drafts in both languages against your style guide and flags anything that reads like a promise you can't keep.

Staff FAQ in two languages

Answers employee questions from your approved internal FAQ only, in English or Spanish, and says so when the answer isn't there.

Minutes in your format

Turns meeting notes into minutes with decisions, owners and due dates, in the template you already use.


Procedure lookup

Finds the relevant section of a policy or procedure and quotes it with the document name, instead of paraphrasing from memory.

MCP: connecting AI to the systems you use

- ✓ **The short version.** A connector lets Claude read one system you choose. Start read-only. Payments and regulated records stay off.

MCP, the [Model Context Protocol](#), is an open standard for connecting AI to other software. A connector built on MCP lets Claude read from one specific system, and act in it only if you allow that, under permissions you set. Think of it as a universal adapter: build the connector once, and any approved AI tool can use it.




How a connector works. Claude reaches a system only through a connector, and only with the permissions you give it. Some systems it never reaches.

Plenty of software has no AI connection today, especially older line-of-business systems. Building the MCP connector for a system that doesn't have one is part of what Masson IT does, and it always starts with the read-only tier.

Build, buy, or leave alone

- ✓ **The short version.** Build small tools that read and draft. Buy systems of record. Leave decisions about clients and money alone.

Building software got much cheaper. Claude can produce a working internal tool in an afternoon. That changes what's worth building, but not what's safe to build. The first question is always what happens if the tool is wrong. Software architect Martin Fowler puts the second one simply: [is it a differentiator?](#) If not, buy it.



The build-or-buy test. Answer the questions in order. The first "yes" decides.

BUILD WITH CLAUDE

Tools that read, sort, draft and fill in. A person checks before anything is saved.

- Memos, minutes and proposals
- Summaries of policies and procedures
- Checklists that flag a missing document
- English and Spanish drafting
- Internal trackers and reports from data you already export
- Inbox search and form-to-screen entry, once the tool is cleared for client information

BUY FROM A SPECIALIST

Systems of record, and anything an auditor or regulator tests. You oversee the vendor.

- Accounting, payroll and payments
- Your industry system of record (medical records, property management, core banking)
- Security and backup tools

Many of these vendors are adding AI features. Use theirs, after review.

LEAVE ALONE

Decisions about people's money, health or eligibility made by a tool you built yourself

- Anything that decides for a client without a person
- Writing into financial or regulated records
- Anything that sends money or messages on its own

LOWER RISK

HIGHER RISK

1

Draft a memo or client notice

BUILD

2

Summarize the employee handbook

BUILD

3

Search your inbox and draft replies

REVIEW FIRST

4

Type paper forms into a draft

REVIEW FIRST

5

Run accounting or payroll

BUY

6

Approve, pay or move money on its own

LEAVE ALONE

Everyday tasks on the risk ladder. A work inbox and most paper forms hold client information, so they get reviewed before Claude touches them.

Six questions before you build

If it's wrong, who is affected?

If the answer is a client or your books, don't build it in-house.

Does it only read and draft, or does it change records?

Tools you build read, sort and draft. A person saves.

Does it touch client information?

Then a privacy and vendor review comes first, under business terms.

Is it what makes you different?

Work that is yours alone can be worth building. What every company does the same way is worth buying.

Who owns it next year?

Name who maintains it, who reviews it, and when it gets retired. Cheap to build is not cheap to own.

Can someone on your team explain what it does?

If not, it doesn't go into daily use.

Buying moves the work, not the responsibility. A bought system still needs someone watching it.

WHERE AN ENGINEER COMES IN


Small tools for your own desk need judgment, not engineering. Once a tool connects to a company system, handles client information, or other people depend on it every day, an engineer should decide whether and how it's built: access, logging, testing, and what happens when it fails. That's where Masson IT's role changes from teaching to building.

08 · WEBSITE

Your website and AI answers

- ✓ **The short version.** AI assistants answer from your website, so give them clear answers in both languages.

More people now ask an AI assistant instead of searching, and they get one answer with a few cited sources. Your website is what those assistants read. Five questions worth asking about yours:



Both languages

Does every page and answer exist in English and Spanish, with the same detail?

Clear answers

Are your most common questions (hours, location, how to get started, what it costs to begin) answered in plain sentences an assistant can quote?

Business details as structured data

Are your address, hours and phone marked up so search engines and AI assistants read them correctly?

FAQ coverage

Which questions does your front desk hear every week that the site doesn't answer yet?

Accessibility

Does it work well on phones and with screen readers?

Skills and connectors help here too. With a read-only connector to your site and a brand-voice skill, your team can draft updates in both languages in minutes. A person still reviews and publishes.

09 · PATH

A path, one step at a time

✓ **The short version.** Learn, set the rules, try one thing, measure it, then grow.

No customer data · 1-5

After review · 6

repeat · 7

1

Pre-work

Your team takes the free AI Fluency for Small Businesses course and earns the certificate.

2

Foundation

Choose the approved tool, set up sign-on and admin controls, and write the one-page AI use policy and data rule.

3

Workshop

Live sessions by team, on real tasks with public or internal information. You leave with a list of where AI could help.

4

Pilot

One workflow, no client data, measured before and after so the result is clear.

5

Skills

Build the playbooks that came out of the workshop a few at a time, and keep the ones people use.

6

Connectors

Read-only connectors first. More only after review.

7

Regular demos

On a set schedule, a short demo of what was built, what stuck and what's next.

The path. Steps 1 to 5 use no customer data. Connectors come after review. Demos repeat on a set schedule.

We don't promise every experiment lands. We promise a steady stream of attempts, a regular demo, and a growing list of things that stuck.

10 • RESOURCES

Free resources

[AI Fluency for Small Businesses](#)

The 4 Ds for business teams, with certificate

[Claude 101](#)

A non-technical starting point

[Introduction to agent skills](#)

How skills are built and used

[agentskills.io](#)

The open skills standard

[AI Fluency: Framework & Foundations](#)

The base course on the 4D framework

[Introduction to Claude Cowork](#)

Working with Claude on real files

[Introduction to Model Context Protocol](#)

How connectors work

[modelcontextprotocol.io](#)

The MCP standard

Want this run for your team?

I teach this guide on site, in English or Spanish, then build the skills and connectors your team needs.

CALL OR TEXT **786 909 3661** • [massonit.com](#)



Masson IT

MASSONIT.COM • MIAMI

Claude, Claude Cowork and Claude Code are Anthropic products. Masson IT is an independent firm and is not affiliated with Anthropic. Course and plan details are as Anthropic published them in fall 2026; confirm before relying on them. This guide is educational and is not legal or compliance advice.