The Complete Claude Guide

From AI fundamentals to full AI workflow

Part 1: AI Foundations

AI fundamentals, prompting, thinking partner, toolkit, cloud vs local

Part 2: AI Governance & Safety

Policy, PII, confidentiality, human-in-the-loop, starting small

Part 3: Stage 1 — Claude Desktop & Web

Account, privacy, integrations, skills, Excel

Part 4: Stage 2 — Claude Code (CLI)

Terminal agent, plugins, IDE, and how it all connects

Resources

GitHub repos, creators, QR codes, links

Plex Consulting

March 2026

For training session attendees

1 Types of Artificial Intelligence

Type	Definition	Status
Narrow AI (ANI)	AI that excels at one specific task — image recognition, language generation, playing chess	This is ALL current AI, including Claude, ChatGPT, Gemini
General AI (AGI)	AI that can learn and perform any intellectual task a human can do	Does not exist yet. Estimated 5–20+ years away (hotly debated)
Super AI (ASI)	AI that surpasses human intelligence across every domain	Theoretical. The focus of AI safety research and debate

Key takeaway: Every AI tool you use today is Narrow AI. It is extraordinarily capable within its domain but has no consciousness, desires, or general understanding. Claude is a language model — it processes text exceptionally well, but it does not "think" the way humans do.

2 Generative vs Non-Generative AI

Generative AI	Non-Generative AI
Creates **new content** — text, images, code, music	Analyses, classifies, or predicts from **existing data**
Claude, ChatGPT, Midjourney, Suno	Spam filters, recommendation engines, fraud detection
Trained on vast datasets, generates probabilistic outputs	Trained on labelled data, outputs decisions or categories
Requires careful prompting to get best results	Requires good training data and feature engineering

3 Open vs Closed Models

Closed-Source (Proprietary)	Open-Source / Open-Weight
Claude (Anthropic), GPT-4 (OpenAI), Gemini (Google)	Llama 3 (Meta), Mistral, Phi-3 (Microsoft), Qwen (Alibaba)
Accessed via API or subscription — you don't see the weights	Weights published — you can download, modify, host yourself
Best performance, easiest to use, vendor lock-in	Run locally, full control, no data leaves your network
Data goes to provider's servers (privacy policies apply)	100% private if self-hosted — nothing phones home

For most businesses: Start with a closed model (Claude Pro) for capability. Add local models later for privacy-sensitive workloads. You don't have to choose one or the other.

4 Key AI Techniques You'll Hear About

Technique	What It Is	Business Example
RAG (Retrieval-Augmented Generation)	AI searches your documents first, then generates answers grounded in your data	Ask Claude a question and it pulls from your company's knowledge base before answering
Fine-Tuning	Training a model on your specific data to specialise its behaviour	Training a model on your legal contracts so it writes in your style
Embeddings	Converting text into numerical vectors so AI can search by meaning, not just keywords	Semantic search: "show me complaints about late delivery" finds relevant tickets even if they don't use those words
Agents	AI that can plan, use tools, and take actions autonomously	Claude Code reads your codebase, writes tests, runs them, and fixes failures — without manual steps

5 How LLMs Work (Simplified)

1. **Training:** The model reads billions of pages of text (books, websites, code, papers) and learns patterns — grammar, facts, reasoning chains, code syntax
2. **Tokenisation:** Your input is split into "tokens" (roughly word fragments). "Understanding" = ["Under", "stand", "ing"]
3. **Prediction:** The model predicts the most probable next token, one at a time, based on everything before it
4. **Context window:** The model can only "see" a fixed amount of text at once (Claude: 200K tokens ≈ 500 pages)
5. **No memory between sessions:** Each conversation starts fresh unless you use Projects, Memory, or CLAUDE.md

> **Common misconception:** LLMs don't "know" things the way humans do. They predict statistically likely continuations. This is why they can be confidently wrong (hallucinations) and why grounding them with your data (RAG, Projects) is essential.

1 The Prompt Framework: R-C-T-F

Element	What It Does	Example
Role	Tell Claude who to be	"You are a senior financial analyst with 15 years in NZ markets"
Context	Give background information	"We're a $50M B2B distributor. Q2 revenue was down 8%."
Task	State exactly what you want	"Analyse the top 3 causes and recommend corrective actions"
Format	Specify the output format	"Present as a 1-page executive summary with bullet points"

> **You don't need all four every time.** For quick questions, just Task is fine. For important work, use all four. The more context you give, the better the output.

2 Iteration Techniques

- **Refine:** "Make it more concise" / "Add more detail to section 2" / "Change the tone to be more formal"
- **Redirect:** "Actually, approach this from the customer's perspective instead"
- **Expand:** "Now take point 3 and develop it into a full paragraph with evidence"
- **Challenge:** "What are the weaknesses in this argument?" / "Play devil's advocate"
- **Constrain:** "Rewrite this in exactly 100 words" / "Use only data from the attached report"

The best outputs come from 2–3 rounds of iteration, not one perfect prompt.

3 Chain-of-Thought & Few-Shot

Chain-of-thought: Ask Claude to reason step by step before giving an answer.

```
"Think through this step by step before giving your final answer:
Should we open a warehouse in Christchurch given our current logistics costs?"
```

Few-shot: Give Claude examples of what you want, then ask for the same pattern.

```
"Here are two examples of how I write customer emails:

Example 1: [paste email]
Example 2: [paste email]

Now write a similar email for a customer whose order arrived damaged."
```

4 Good vs Bad Prompts

Bad Prompt	Good Prompt
"Write a marketing email"	"Write a 200-word marketing email for our Q2 office supplies sale targeting NZ schools. Tone: professional but warm. Include a clear CTA to our landing page."
"Summarise this document"	"Summarise this 40-page report in 5 bullet points. Focus on financial risks and recommended actions. Audience: the board."
"Help me with strategy"	"You are a strategy consultant. We're a $50M distributor losing market share to online competitors. List 5 strategic options with pros, cons, and estimated cost to implement."
"Fix this code"	"This Python function returns None when it should return a list. Here's the function and the error trace. Explain the bug, fix it, and add a unit test."

5 Common Pitfalls

Pitfall	Why It Happens	Fix
Hallucinations	Model generates plausible but false information	Ask "cite your sources" / use web search / upload reference docs
Too generic	Not enough context or constraints	Add role, specifics, format, examples
Too long	No length constraint specified	"Maximum 200 words" / "3 bullet points only"
Wrong tone	No tone guidance given	"Write in a formal/casual/technical tone for [audience]"
Ignores instructions	Too many competing instructions	Simplify. Put the most important instruction last. Use numbered steps.

Thinking Partner Prompts

10 READY-TO-USE PROMPTS

Plex Consulting

Use these prompts to turn Claude into a strategic thinking partner. Paste any of them directly into a conversation. Adapted from **@natan_mohart**.

1. Challenge My Thinking

"I'm considering [decision/strategy]. Act as a sharp, experienced advisor. Challenge every assumption I've made. Point out what I might be wrong about. Be direct — I want intellectual honesty, not validation."

2. Reframe the Problem

"Here's my current problem: [describe it]. I suspect I'm framing it wrong. Give me 5 completely different ways to frame this problem. For each reframe, explain what solution it would lead to."

3. Surface Unstated Questions

"I'm about to [decision]. Before I commit, what are the 10 questions I should be asking but haven't thought of? Focus on blind spots, second-order effects, and things that could go wrong in 6–12 months."

4. Find Second-Order Effects

"If we [action], what are the second and third-order effects? Think beyond the obvious first-order outcomes. Consider impacts on team morale, customer perception, competitor response, operational complexity, and cash flow."

5. Reveal Hidden Variables

"I'm evaluating [situation] and I've identified these factors: [list them]. What important variables am I not considering? What data should I be looking at that I'm probably ignoring?"

6. Extract Principles

"Here are 5 past decisions I made that worked well: [list them]. And 3 that didn't: [list them]. What underlying principles or patterns can you extract? Give me a personal decision-making framework based on my own history."

7. Reverse-Engineer Instinct

"My gut says [instinct about a decision]. I can't fully articulate why. Help me reverse-engineer my intuition. What might I be pattern-matching on? What past experience could be driving this feeling? Is my gut likely right here?"

8. Find Patterns

"Here's a list of [data/events/feedback]: [paste them]. What patterns do you see that I might be missing? Group them, rank them by significance, and tell me what the patterns suggest about what's really happening."

9. Think in Reverse

"I want to achieve [goal] by [date]. Work backwards from the end state. What needs to be true 1 month before? 3 months before? 6 months before? What's the critical path and where are the bottlenecks?"

10. Translate Intuition

"I have a vague sense that [fuzzy idea]. I can't quite put it into words. Help me articulate this clearly. Ask me clarifying questions, then write a crisp 2-paragraph version of what I'm trying to say."

1 Tier 1: Daily Drivers (~$60/month)

These are the tools you'll use every single day. Start here.

Tool	Cost	What It Does
Claude Pro	$20/mo (USD)	Your primary AI assistant. Chat, research, documents, code, Excel add-in. Opus 4.6 is the best model available anywhere. Includes Claude Code and web connectors.
Cursor Pro	$20/mo (USD)	AI-native code editor. Uses Claude under the hood. Tab-completion, inline editing, multi-file refactoring. The best IDE for AI-assisted development.
Perplexity Pro	$20/mo (USD)	AI-powered search engine with citations. When you need factual, up-to-date, sourced answers rather than generated text. Best for research and fact-checking.

2 Tier 2: Build & Automate (Free–$30/month)

Tool	Cost	What It Does
Figma + AI	Free–$15/mo	Design tool with AI features. Figma Make generates designs from prompts. MCP connector lets Claude interact with your Figma files directly.
Git + Cloudflare	Free	Version control (Git/GitHub) + free hosting (Cloudflare Pages). Deploy websites and tools for free. Claude Code manages git natively.
n8n	Free (self-hosted)	Visual workflow automation. Connect AI to email, databases, APIs, spreadsheets. Trigger Claude workflows on schedules, webhooks, or events.

3 Tier 3: Infrastructure ($0–$100+/month)

Tool	Cost	What It Does
Ollama	Free	Run open-source LLMs locally (Llama 3, Mistral, Phi-3). 100% private. Needs a decent GPU (8GB+ VRAM recommended). Great for experimentation.
AWS Bedrock	Pay-per-use	Run Claude, Llama, Mistral in the cloud via API. Enterprise-grade. Good when you need Claude API access beyond the Pro subscription.
JARVIS	Custom build	Multi-agent AI platform (what Plex runs). 254 specialist agents, 25 MCP servers, knowledge base, automation. The "endgame" for AI-native businesses.

4 Tier 4: Occasional Use (Free–$20/month)

Tool	Cost	When to Use
ChatGPT	Free–$20/mo	Image generation (DALL-E), voice mode, mobile app. Second opinion on complex tasks. GPT-4o is good but Claude Opus is better for reasoning and writing.
NotebookLM	Free (Google)	Upload documents and get an AI that only answers from those docs. Creates AI-generated podcast discussions from your content. Great for study/review.
Gemini	Free–$20/mo	Google's AI. Excellent for Google Workspace integration. 1M+ token context. Good for analysing very large documents.

$ Monthly Cost Summary

Tier	Tools	Monthly (NZD approx.)
Tier 1: Daily Drivers	Claude + Cursor + Perplexity	~$105
Tier 2: Build & Automate	Figma + Git + n8n	$0–$25
Tier 3: Infrastructure	Ollama + Bedrock	$0–$170+
Tier 4: Occasional	ChatGPT + NotebookLM + Gemini	$0–$35
Realistic Total (Tier 1 + selective Tier 2–4)		**~$105–$210/mo**

Start with Tier 1 only. Claude Pro at $20/month is the single best investment in AI productivity. Add tools as you hit specific needs, not before. Most people never need Tier 3.

Cloud vs Local AI

THE PRIVACY AND COST TRADE-OFF

Plex Consulting

Cloud AI — Pros & Cons

+/–	Detail
+	**Best models:** Claude Opus 4.6, GPT-4o, Gemini Ultra — far superior to anything you can run locally
+	**Zero setup:** Sign up, start using. No hardware, no maintenance, no GPU required
+	**Always updated:** You get model improvements automatically. No need to download new versions
+	**Massive context:** Claude handles 200K tokens (500+ pages) — local models typically max out at 8K–32K
–	**Data leaves your network:** Your prompts and data are sent to the provider's servers
–	**Ongoing cost:** $20–$200+/month depending on usage and tier
–	**Vendor dependency:** If the provider changes pricing, policies, or goes down, you're affected
–	**Compliance concerns:** Some industries (legal, health, defence) may restrict cloud AI use

Local AI — Pros & Cons

+/−	Detail
+	**Total privacy:** Nothing leaves your machine. Zero data exposure. Full compliance
+	**No ongoing cost:** After hardware purchase, running costs are just electricity
+	**No internet required:** Works offline, on planes, in secure facilities
+	**Full control:** Choose models, fine-tune, customise, no terms of service restrictions
–	**Lower quality:** Best local models (Llama 3 70B) are good but noticeably behind Opus/GPT-4o
–	**Hardware cost:** Need a GPU with sufficient VRAM. Quality models need 24GB+ VRAM
–	**Technical setup:** Requires command-line comfort, driver management, model selection
–	**Smaller context:** Most local models handle 4K–32K tokens vs cloud's 200K+

Hardware Costs (NZD, March 2026)

Setup	Hardware	Cost (NZD)	What You Can Run
Entry	Used gaming PC with RTX 3060 12GB	$800–$1,200	7B–13B models (Llama 3 8B, Mistral 7B). Good for simple tasks, drafting, summarisation
Mid	RTX 4070 Ti Super 16GB or RTX 3090 24GB	$1,200–$2,500	13B–34B models (Llama 3 70B quantised). Solid general-purpose local AI
Pro	RTX 4090 24GB or dual GPU	$3,500–$6,000	70B models at full quality. Comparable to cloud for many tasks
Server	Apple Mac Studio M2/M4 Ultra 192GB	$8,000–$14,000	Full 70B+ models, multiple concurrent users. The "always-on" local AI server

The Hybrid Answer

Most businesses should use both. The optimal setup:

- **Cloud (Claude Pro):** Complex reasoning, research, document creation, coding — where quality matters most
- **Local (Ollama):** Sensitive data processing, internal document analysis, embeddings, classification — where privacy matters most

> **The business case is simple:** Claude Pro costs $20/month. A local GPU setup starts at $800+. You need 40+ months of Claude Pro to break even on hardware alone — and Claude Pro is better. Start with cloud. Add local only when you have a specific privacy or compliance requirement that demands it.

! Why This Matters

AI is not magic — it is a tool. Like any tool, it needs rules. The organisations that move fastest with AI are the ones that set clear boundaries first, not the ones that figure it out after something goes wrong.

The biggest risk is not that AI will "go rogue". It's that someone pastes confidential data into the wrong tool, or trusts an AI answer without checking, or automates a process that should have had a human checkpoint.

1 Confidential Information & PII Policy

Every organisation needs a clear AI data policy. Before anyone uses AI, they need to know:

Question	Your Policy Should Answer
What data can go in?	Define categories: public (OK), internal (maybe), confidential (restricted), PII (never without controls). Be specific — "don't put sensitive stuff in" is not a policy.
Which tools are approved?	Name the specific AI tools your org approves (e.g. "Claude Pro with training opt-out"). Ban unapproved tools. Free-tier AI tools often train on your data.
Where does data go?	Cloud AI sends data to a server. Know which country, which provider, what their retention policy is. Claude Pro/Team/Enterprise: data is NOT used for training by default.
What about PII?	Names, emails, phone numbers, health info, financial details — never paste these into AI unless you have a compliant enterprise plan with data processing agreements.
Who is accountable?	AI outputs need an owner. If Claude drafts a client email, a human is responsible for what gets sent. Always.

> **The #1 mistake:** People paste customer spreadsheets, employee records, or financial data into free AI tools. Those tools may use your data for training. Use Enterprise/Team plans or local models for anything sensitive. Free tier = your data may be used.

2 Data Classification for AI Use

Level	Examples	AI Tool Allowed	Rules
PUBLIC	Marketing copy, public web content, press releases	Any AI tool (including free tier)	No restrictions
INTERNAL	Meeting notes, internal reports, strategy drafts, process docs	Approved paid tools only (training opt-out ON)	Review output before sharing externally
CONFIDENTIAL	Financial results, contracts, board papers, IP, trade secrets	Enterprise plan OR local AI only	Human review mandatory. No copy-paste of raw data. Anonymise first if possible.
RESTRICTED / PII	Customer names+addresses, employee records, health data, credit cards	Local AI only (or DPA-compliant enterprise)	Never in cloud AI without a Data Processing Agreement. Anonymise or use synthetic data.

3 Enforcement — Making Policy Stick

A policy nobody follows is worse than no policy — it gives false confidence. Here's how to enforce:

Mechanism	How It Works
Approved tool list	Publish a list of approved AI tools. Block unapproved ones at the network/IT level if possible. Review quarterly as new tools emerge.
Training opt-out verification	For every approved tool, confirm the "do not train on my data" setting is ON. Screenshot it. Check it monthly. Claude Pro: Settings → Privacy → "Improve Claude" = OFF.
Prompt hygiene training	Train staff: "Before you paste, ask: would I email this to a stranger?" If no, don't paste it into AI. 30-minute session, repeat annually.
Spot audits	Randomly review AI usage logs (Enterprise plans provide these). Not to punish — to find gaps in training and policy.
Incident process	If someone accidentally puts PII into AI: who do they tell? What happens next? Make it safe to report. No blame, just fix.
Acceptable Use Agreement	Add an AI clause to your existing AUA/IT policy. Staff sign it. Simple: "I will only use approved AI tools and will not input confidential data without authorisation."

The Fear: "AI Will Go Out of Control"

This is the most common concern. The answer is simple: **AI only does what you let it do.** Every AI tool has control points where humans decide what happens next. The key is knowing where those points are and never bypassing them.

Think of it like giving a new employee their first week of tasks. You don't hand them the company credit card and say "do whatever you think is best." You give them small tasks, check the output, and gradually increase responsibility as trust is earned.

1 The Confidence Ladder — Start Here

Step	What You Do	Human Control Point	When to Move Up
1	**Read-only** AI summarises, explains, answers questions	Human reads the output and decides if it's useful. AI changes nothing.	When you trust the quality of answers (1–2 weeks)
2	**Draft mode** AI writes drafts (emails, reports, proposals)	Human reviews and edits every draft before it goes anywhere. AI is a ghostwriter.	When you're editing less than 20% of drafts (2–4 weeks)
3	**Assisted decisions** AI analyses data, recommends options	Human chooses which recommendation to act on. AI never decides alone.	When recommendations align with your judgement 80%+ of the time
4	**Supervised automation** AI executes tasks with approval gates	AI proposes action, waits for human "yes" before executing. Like Plan Mode in Claude Code.	When approval rate is 95%+ and exceptions are well-handled
5	**Autonomous (bounded)** AI runs within strict guardrails	Human sets the boundaries. AI operates within them. Alerts on anything outside scope.	Only for well-understood, low-risk, repeatable tasks

> **Key principle:** You earn the right to move up each step. Don't jump to step 5 on day one. Most organisations should spend 1–3 months at steps 1–2 before considering automation.

2 Where to Put the Human Checkpoints

Scenario	Human Checkpoint	Why
AI drafts a client email	Human reads and sends manually	Tone, accuracy, relationship context that AI doesn't have
AI analyses financial data	Human verifies numbers against source	AI can hallucinate figures. Always cross-check numbers.
AI writes code	Human reviews diff before committing	Security, logic errors, unintended side effects
AI summarises a meeting	Human checks key decisions are captured correctly	AI may miss nuance, sarcasm, or unspoken agreements
AI recommends a strategy	Human evaluates against business context	AI lacks institutional knowledge, politics, and timing awareness
AI generates a report for the board	Human fact-checks every claim and number	Your name is on it, not the AI's
AI automates a workflow (n8n, etc.)	Human approves the first 10 runs manually	Catch edge cases before they become automated mistakes

3 Built-in Safety Features in Claude

Claude has specific features designed to keep humans in control:

Feature	What It Does	Where to Find It
Plan Mode	Claude proposes changes and waits for your approval before executing anything	Claude Code: `Shift+Tab` or `/plan`
Permission System	You control which tools Claude can use (file access, terminal, web). Deny anything you're not comfortable with.	Claude Code: `/permissions`
Training Opt-Out	Your data is NOT used to train models (Pro/Team/Enterprise)	Settings → Privacy → "Improve Claude" = OFF
Memory Controls	Review, edit, or delete anything Claude remembers about you	Settings → Memory (view/delete individual items)
Incognito Mode	Conversations are not stored and memory is disabled for that session	Start a new chat with incognito toggle
Hooks & Guardrails	Custom rules that block or warn before specific actions (e.g., block .env file access)	Claude Code: `/hooks` or settings.json
Audit Logs	Full history of what Claude did in each session (Enterprise/Team)	Admin console → Usage & Audit

★ The 5-Point Safety Checklist

Before rolling out AI to your team, confirm these five things:

1. **Data policy exists and is signed** — everyone knows what data can and cannot go into AI
2. **Training opt-out is verified** — confirmed and screenshotted for every approved tool
3. **Human review is mandatory for external outputs** — nothing AI-generated goes to a client, regulator, or the public without a human checking it
4. **Start at Step 1 (read-only)** — resist the urge to automate on day one. Build confidence first.
5. **Incident process is in place** — if something goes wrong, people know who to tell and it's safe to report

> **The bottom line:** AI is a power tool, not a self-driving car. You hold the steering wheel. The organisations that succeed with AI are the ones that stay in the driver's seat while letting AI handle the heavy lifting. Start small. Check everything. Gradually trust more as you earn confidence.

Stage 1 — Claude Desktop & Web

Everything you need with zero technical setup. 10 minutes to fully operational.

1 Sign Up for Claude Pro ($20/month)

1. Go to **claude.ai** and click **Sign Up**
2. Create an account with your email (or Google/Apple sign-in)
3. Click your name (bottom-left) → **Upgrade to Pro**
4. Enter payment details — you're now on Claude Pro

> **What Pro gets you:** Opus 4.6 (best model), Projects, Memory, Research, Skills, Claude Code, Excel add-in, web search, file creation, and all connectors. Free tier is very limited — Pro is where it starts.

2 Download the Desktop App

1. Go to **claude.ai/download**
2. Download for Windows (or Mac)
3. Install and sign in with your account

The desktop app gives you the same features as the web, plus one-click MCP server installs (Desktop Extensions) and better file access.

3 Set Up Privacy & Data Controls

This is important. By default, Pro plan data is not used for training. But check these settings:

1. Click your **name** (bottom-left) → **Settings**
2. Go to **Privacy & Data** (or **Data Controls**)
3. Confirm these settings:

Setting	Recommended	What It Means
Improve Claude	OFF	Your conversations are NOT used to train future models
Chat History	ON	You can see and search past conversations
Memory	ON	Claude remembers your preferences across sessions

> **Enterprise note:** Team and Enterprise plans have admin-controlled data policies, SSO, audit logs, and Zero Data Retention (ZDR) options. If your company requires these, speak to your IT team about an Enterprise plan.

4 Configure Memory

Memory lets Claude learn your preferences, role, projects, and communication style over time.

1. **Settings** → **Capabilities** → ensure **Memory** is ON
2. In any conversation, tell Claude things to remember:

```
"Remember that I'm the CEO of Acme Corp. I prefer bullet points over paragraphs.
My team uses Slack for comms and Monday.com for project tracking.
Always use NZ English spelling."
```

Claude auto-summarises conversations every 24 hours and builds a profile. You can review what it remembers anytime:

- **Settings** → **Memory** → view, edit, or delete saved memories

> **Tip:** The more context you give Claude upfront, the better it performs. Spend 5 minutes telling it who you are, what you do, and how you like to work. It pays dividends in every future session.

5 Create Your First Project

Projects are persistent workspaces. Each one has its own uploaded documents, custom instructions, and conversation history.

1. Click **Projects** in the left sidebar
2. Click **+ New Project**
3. Give it a name (e.g. "Q2 Strategy", "Marketing Copy", "Board Prep")
4. Add **Custom Instructions** — tell Claude how to behave in this project:

```
"You are helping me prepare board papers for Acme Corp.
Always use formal business English. Reference NZ market data where relevant.
Structure documents with executive summary first, then detail sections.
Financial figures in NZD. Use FY26 as the current financial year."
```

1. **Upload documents** — drag in PDFs, spreadsheets, reports. Claude reads them all (200K context = ~500 pages)
2. Start chatting — Claude uses everything you've uploaded as context

> **Project = your knowledge base.** Upload your company's strategy doc, financial reports, brand guide, org chart — whatever Claude needs to give you informed answers. Every conversation inside that project has access to all of it.

6 Connect Your Tools (Integrations)

Integrations let Claude search and pull data from your existing tools — no copy-pasting needed.

In the Claude Desktop App:

1. Click **Settings** (gear icon) → **Extensions** or **Integrations**
2. Browse available connectors and click **Connect** on the ones you use

On claude.ai web:

1. Visit **claude.ai/directory** to browse 75+ available connectors
2. Click **Connect** and authorise with your account

Integration	What Claude Can Do
Google Workspace	Search Gmail, read Calendar, access Google Docs & Drive files
Slack	Search messages, read channels, find conversations
Notion	Search pages, read databases, access workspace content
Jira / Confluence	Read tickets, search documentation, access project data
Asana	Read tasks, search projects, access team workspaces
Zapier	Connect to 1,000+ apps through Zapier's automation platform
Linear	Search issues, read project status, access team boards

> **How it works:** Once connected, you can ask Claude things like "check my Gmail for the latest email from Sarah" or "what's on my calendar tomorrow" — Claude pulls the live data directly. Your credentials are stored securely in your Anthropic account.

7 Explore Skills

Skills are pre-built capability packs that make Claude instantly better at specific tasks.

Built-in Skills (already available on Pro):

Skill	What It Does
Excel (xlsx)	Creates formatted spreadsheets with formulas, charts, multiple tabs
Word (docx)	Creates professional Word documents with formatting and styles
PowerPoint (pptx)	Creates slide decks with layouts, charts, and speaker notes
PDF	Creates formatted PDF documents
Research	Deep multi-agent research across web and connected services

To use a skill, just ask Claude naturally:

```
"Create a PowerPoint presentation summarising our Q2 results"
"Build an Excel spreadsheet tracking our marketing budget with formulas"
"Research the NZ construction market and write a 3-page report as a PDF"
```

Claude automatically activates the right skill based on your request. The output appears as a downloadable file.

8 Install Claude in Excel

The official Microsoft Excel add-in puts Claude inside your spreadsheets.

1. Open **Excel** (2016 or later, or Excel Web)
2. Go to **Insert** → **Get Add-ins** (or **Office Add-ins**)
3. Search for **"Claude"**
4. Click **Add** → sign in with your Anthropic account

What it can do:

- Read your actual formulas, cell references, and tab structure — not just the visible text
- Explain complex formulas in plain English
- Debug formula errors and suggest fixes
- Write new formulas based on what you describe
- Analyse data patterns across your workbook
- Generate charts and summaries from your data

> **Example:** Select a range of cells, open the Claude panel, and ask "What's wrong with this VLOOKUP?" or "Create a pivot summary of this data by region and quarter". Claude sees the actual spreadsheet structure, not a screenshot.

9 Use Artifacts for Live Documents

Artifacts are interactive outputs that appear in a dedicated panel alongside your chat.

- **Documents** — reports, proposals, meeting notes that you can edit live
- **Calculators** — interactive tools (ROI calculator, pricing model, budget tracker)
- **Dashboards** — visual data displays with charts
- **Web apps** — functional mini-applications (forms, trackers, tools)
- **Diagrams** — org charts, process flows, architecture diagrams

Ask Claude to "create an artifact" or just describe what you need — it'll create one automatically when appropriate.

```
"Build me an interactive ROI calculator for our SaaS product"
"Create a project timeline diagram for our office move"
"Make a comparison table of these 5 vendors as a shareable document"
```

Stage 1 Complete

You now have:

- Claude Pro with privacy configured
- Memory learning your preferences
- Projects as persistent knowledge bases
- Integrations pulling data from your tools
- Skills creating professional documents
- Claude inside Excel
- Artifacts for interactive outputs

This is enough for most people. Stage 2 is for those who want to go deeper with code, automation, and the terminal agent.

Stage 2 — Claude Code (CLI)

The terminal-based AI agent. Writes code, manages files, runs commands, and connects to everything.

What is Claude Code?

Claude Code is an **AI agent that lives in your terminal**. Unlike the chat interface, it can:

- **Read and write files** on your computer directly
- **Run terminal commands** (build, test, deploy, git)
- **Edit code** with full understanding of your project structure
- **Search your codebase** instantly across thousands of files
- **Spawn sub-agents** for parallel work
- **Connect to MCP servers** for external tool access

It's included free with your Claude Pro subscription. Same account, same models.

1 Install Claude Code

Automated Setup Script

Download: `setup-claude-code.ps1`

Run:

```
powershell -ExecutionPolicy Bypass -File setup-claude-code.ps1
```

This installs Node.js, Claude Code, plugins, and the VS Code extension in one go. Interactive — you pick what you want.

Option A: Use the PowerShell setup script (recommended — see above)

Option B: Manual install

```
# Install Node.js first from nodejs.org, then:
npm install -g @anthropic-ai/claude-code
```

Option C: VS Code only

1. Open VS Code
2. Press **Ctrl+Shift+X** (Extensions)
3. Search **"Claude Code"**
4. Click **Install**

2 First Run & Authentication

1. Open a terminal in any project folder
2. Type `claude` and press Enter
3. A browser window opens — sign in with your Anthropic account
4. Return to the terminal — you're connected

First thing to do:

```
/init
```

This scans your project and generates a `CLAUDE.md` file — a set of instructions that Claude reads at the start of every session. Think of it as your project's briefing document for the AI.

3 Memory in Claude Code

Claude Code has **two memory systems**:

System	Who Writes It	What It Contains
CLAUDE.md	You	Project instructions, coding standards, build commands, team conventions
Auto Memory	Claude	Things it learns during sessions — your patterns, preferences, corrections

Both are loaded automatically at the start of every session. Claude Code's memory is **file-based and local** — it lives on your machine, not in the cloud.

Key commands:

Command	What It Does
`/memory`	View and manage what Claude remembers
`/compact`	Compress conversation to free up context space
`/context`	Check how much context window you've used
`/model`	Switch between Opus, Sonnet, and Haiku
`/effort`	Set thinking depth (more effort = deeper reasoning)
`/plan`	Safe mode — Claude proposes changes, waits for approval

4 Plugins & Integrations

Plugins extend Claude Code with new capabilities. Install from the official marketplace:

```
# Core plugins (recommended for everyone)
/plugin install code-review@claude-code-marketplace
/plugin install security-guidance@claude-code-marketplace
/plugin install frontend-design@claude-code-marketplace

# Language intelligence (pick your language)
/plugin install typescript-lsp@claude-code-marketplace
/plugin install pyright-lsp@claude-code-marketplace
/plugin install csharp-lsp@claude-code-marketplace

# Tool connectors
/plugin install github@claude-code-marketplace
/plugin install slack@claude-code-marketplace
/plugin install context7@claude-code-marketplace
```

Manage plugins anytime with `/plugin`. View installed plugins with `/plugin list`.

5 Built-in Skills

These ship with every Claude Code installation:

Skill	What It Does	Example
`/batch`	Parallel changes across your codebase using sub-agents	"Rename all API endpoints from camelCase to snake_case"
`/claude-api`	Loads Claude API reference for your language	Auto-activates when writing Claude integrations
`/debug`	Troubleshoots your current session	"Why did that last command fail?"
`/loop`	Runs a prompt on a recurring interval	"/loop 5m check if the build passed"
`/simplify`	Reviews recent changes for quality and reuse	Run after writing code to catch issues

How It All Connects

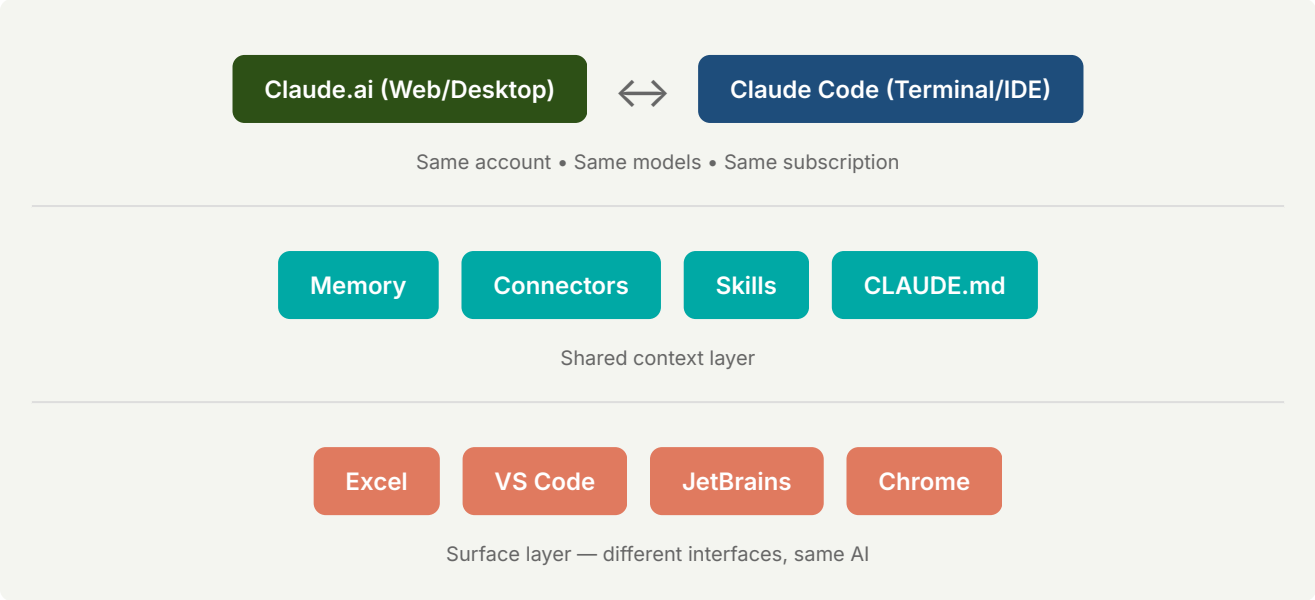

The key insight: It's all one Claude. Your web account, desktop app, CLI, VS Code extension, and Excel add-in all share the same subscription, models, and (where supported) connectors. You don't pay separately for any of these.

What Carries Over	From → To
Account & subscription	Everywhere
Web connectors (Gmail, Slack, Notion)	Claude.ai ↔ Claude Code
Skills (document creation)	Claude.ai ↔ Claude Code
CLAUDE.md instructions	Shared across Claude Code & VS Code
Chat memory (auto-learned)	Claude.ai (cloud-based)
Auto memory (file-based)	Claude Code (local to your machine)

6 Cowork — The Bridge

Claude Cowork is Anthropic's GUI tool that bridges web and code. Think of it as the desktop version of Claude Code with a visual interface.

Feature	Claude.ai	Cowork	Claude Code
Chat interface	✓	✓	✓ (terminal)
File access	Upload only	✓ Local files	✓ Full filesystem
Run commands	×	✓	✓
CLAUDE.md	×	✓	✓
Connectors	✓	✓	✓
Skills	✓	✓	✓
Scheduled tasks	×	✓ (/schedule)	✓ (/loop)
Git integration	×	✓	✓

> **When to use which:**
> **Claude.ai** — thinking, writing, research, quick questions
> **Cowork** — project work that needs file access and command execution, with a GUI
> **Claude Code** — full power, automation, scripting, multi-agent workflows, CI/CD

Quick Reference Card

Action	How
Start Claude Code	`claude` in any terminal
Set up a project	`/init`
Check what Claude remembers	`/memory`
Free up context space	`/compact`
Switch model	`/model` (Opus for quality, Haiku for speed)
Install a plugin	`/plugin install name@claude-code-marketplace`
Review a pull request	`/code-review`
Run something on repeat	`/loop 10m check the deploy status`
Safe mode (ask before acting)	`/plan` or **Shift+Tab** before typing
Check token usage	`/context`
See all commands	`/help`

Resources — Where to Find More

QR codes, GitHub repos, Instagram creators, and community hubs for Claude agents, skills, and tips.

Quick Links (Scan to Open)

All Resources & Downloads
https://aidocs-6ab.pages.dev

Download Desktop App
claude.ai/download

CLI Setup Script
https://aidocs-6ab.pages.dev/tools/

Top GitHub Repos for Agents & Skills

Repository	What It Is
anthropics/skills	Official Anthropic skills — document creation (xlsx, docx, pptx, pdf), frontend design, API reference, and more
anthropics/claude-code	Official Claude Code repo — CLI source, bundled skills, plugin architecture, documentation
anthropics/claude-plugins-official	Official plugin marketplace — 32 Anthropic plugins + 15 partner plugins
modelcontextprotocol/servers	Official MCP servers — filesystem, git, memory, fetch, sequential thinking, and community servers
anthropics/anthropic-cookbook	Code examples and patterns for building with the Claude API
frankbria/ralph-claude-code	Ralph Loop — autonomous development agent loops for long-running coding sessions
thedotmack/claude-mem	Enhanced long-term memory system for Claude Code (was #1 trending on GitHub)

Skills discovery & community:

Site	What You'll Find
agentskills.io	Open standard for agent skills — works across multiple AI tools
operatoros.ai	Agent skills platform by @daviss.dev — /swarm for multi-agent orchestration
claude.ai/directory	Official connector directory — 75+ integrations for Claude Desktop

> **Installing community skills:** Use `npx skills add author/repo@skill-name` in your terminal. For example: `npx skills add anthropics/skills@frontend-design`

★ Top Instagram Creators for Claude Tips

These accounts consistently post practical Claude tutorials, skills, and workflows:

Account	Focus	Why Follow
@mindwired.ai	Claude Code tools & plugins	Best breakdowns of new plugins and GitHub tools. First to cover Ralph Loop, Claude Mem, code review agents.
@ux.amish	AI design skills	Design-focused Claude skills (UI-UX-Pro-Max, accessibility, web design guidelines). Great for designers learning Claude Code.
@os.operator (daviss.dev)	Agent skills & orchestration	Creator of OperatorOS.ai. Multi-agent /swarm patterns, skill architecture, advanced workflows.
@natan_mohart	Thinking partner prompts	Practical prompt frameworks for using Claude as a strategic advisor and thinking partner.
@okaashish	n8n automation	Complete n8n masterclasses — from beginner nodes to complex AI-powered workflows.
@meetavinash	Figma Make + AI	Prompt frameworks for Figma Make — systems architect, visual design, conversion copy.
@analytics_vidhya	AI tutorials & MCP	Technical walkthroughs of MCP setup, Figma integration, Claude Code features.
@theaisolopreneur	AI for business	Practical business use cases, prompt templates, productivity workflows with Claude.
@aiaboratory	AI tools & workflows	Side-by-side tool comparisons, new feature coverage, beginner-friendly tutorials.

> **Tip:** Follow 3–4 of these accounts and spend 5 minutes a day watching their reels. The AI landscape moves fast — Instagram creators often cover new features before the official docs are updated.

☞ Official Anthropic Links

Resource	URL
Claude (sign up / chat)	claude.ai
Claude Code docs	code.claude.com/docs
Connector directory	claude.ai/directory
Pricing & plans	claude.ai/pricing
Desktop app download	claude.ai/download
API documentation	docs.anthropic.com
Anthropic blog	anthropic.com/news
Claude Code GitHub	github.com/anthropics/claude-code