

The Complete Claude Code Implementation Guide

A practical guide for marketing directors, operations leads and hands-on managers ready to put agentic AI to work

ABOUT THE AUTHOR

Ann Stanley

Founder & CTO, Anicca Digital

Co-Founder & CMO, Anicca AI & Insights

I founded **Anicca Digital** in 2007. An award-winning agency known for data-driven ecommerce and digital marketing.

Prior to that I had built a career spanning science, medical research, and manufacturing. My first role in digital marketing came in 2002, so by the time Anicca launched I already had over 20 years of business experience.

I have trained over **1,000 marketers** and spoken to thousands more at conferences and webinars. I am the author of two books, *Integrate* and the *A10 Marketing Framework*, and am currently writing my third, the **C10 AI Adoption Framework**.

Five years ago, before ChatGPT had even launched, Anicca was already investigating how AI could predict the impact of weather on garden furniture sales. When ChatGPT arrived, I wrote my *Integrate* book within a month.

In January 2025, I committed every weekend for **seven months** to intensive AI training. By the end of it I was not just advising on AI, I was building it. I built agents, wrote code, and vibe-coded an entire website. My course project became **secret-agents.ai**, an AI agent marketplace.

By May 2025, Anicca had launched a full range of AI services. Darren Wynn and James Allen were building AI solutions alongside me. We joined forces with Lucia Dello and Paul Hudson to form **Anicca Insights**, developing our Secret Agents marketplace and **Armadello**, an ecommerce analytics app now live on the Shopify App Store.

When I started using **Claude Code** in January 2026, it felt like a seismic shift all over again.

I have since built **over a hundred practical AI agents** covering every aspect of marketing and management.

If you are reading this thinking "I am not technical enough for this" - good.

Neither am I. That is the whole point.

Health Warning

Side effects may include: sore eyes from staring at Claude Code output at 2am, a stiff back from forgetting to move for six hours, and an inability to stop talking about AI to anyone who will listen.

How to Use This Guide

Already sold on AI? Skip to Chapter 2 (page 8) to get set up today.

Need convincing first? Chapter 1 has the data and the strategic context.

Want to see what it can do? Chapter 6 has real examples from our agency.



Contents

CHAPTER 1 - WHY CLAUDE CODE MATTERS NOW

1.1 What Is Claude Code and Why Agentic AI Is Changing How Businesses Operate	1
1.2 The Adoption Gap	2
1.3 Where Are You on the AI Adoption Ladder?	3
1.4 Automated Workflows vs AI Agents	4
1.5 How Humans and AI Work Together	5
1.6 Why Claude Code?	6
1.7 Who Should Implement First	7

CHAPTER 2 - UNDERSTANDING CLAUDE CODE

2.1 The Evolution from Projects to Claude Code	8
2.2 Claude Desktop: The Simplest Way In	9
2.3 Claude in Cursor / VS Code	11
2.4 Choosing Your Setup	12
2.5 The Innovation Edge	13
2.6 Your First Practical Win	14

CHAPTER 3 - GETTING SET UP

3.1 Setting Up: Day 1	16
3.2 Additional Setup, Version Control and Voice Commands	17
3.3 Safety First: How Claude Asks Before It Acts	18
3.4 Software Tools and Costs	19
3.5 Setting Up Cursor or VS Code for Power Users	20

CHAPTER 4 - HOW TO WORK WITH CLAUDE

4.1 The Seven Concepts That Matter	21
4.2 How to Talk to Claude	22
4.3 Your Brand Voice	23
4.4 Controlling Thinking Depth	24
4.5 Your CLAUDE.md Brand Bible	25
4.6 Memory, Handovers and Sessions	27
4.7 Session Management	28

CHAPTER 5 - SKILLS

5.1 A Skill Is Just a Folder	29
5.2 The Seven Levels of Skill Building	30
5.3 Building Your First Skill	31
5.4 The Skill Creator: The Skill That Builds All Others	32
5.5 Reproducible Intelligence: Why Skills Are Your Competitive Moat	33
5.6 The HTML-First Workflow	34
5.7 QC and Critique: Let Claude Check Its Own Work	35
5.8 Your First Week Plan	36
5.9 What to Build Next	37

CHAPTER 6 - BUSINESS USE CASES AND OUTPUTS

6.1 The Benefits Framework	38
6.2 The Full Department Map	39
6.3 Marketing and Content	40
6.4 Documents, Reports and Knowledge	41
6.5 What Types of Content Can Claude Code Produce?	42
6.6 Real Examples: Documents We Have Built	43
6.7 Sales and New Business	44
6.8 Research and Strategy	45
6.9 Finance and Administration	45
6.10 Advanced: Web Apps, Portals and Dashboards	46
6.11 APIs, MCP and Credentials	47
6.12 From Individual Tasks to Automated Workflows	48
6.13 Building Web Applications with a Database	49
6.14 Building Real Applications Without a Developer	50
6.15 Real Examples: Applications We Have Built	51

Contents (continued)

CHAPTER 7 - IMPLEMENTATION AND GOVERNANCE

7.1 Crawl, Walk, Run	52
7.2 What Does It Cost Your Organisation?	53
7.3 Your First Month Health Check	54
7.4 Ambassadors and Adoption Styles	55
7.5 Getting Your Team on Board	56
7.6 Sharing What Claude Builds	57
7.7 Governance	58
7.8 Running Your First AI Discovery Session	59
7.9 The C10 AI Adoption Framework	60

CHAPTER 8 - FROM PERSONAL AI TO COMPANY AI

8.1 From individual tasks to whole systems	61
8.2 Why personal AI does not scale on its own	62
8.3 Two ways businesses arrive at company-wide AI	63
8.4 AI departmental agent products on the market	64
8.5 What an AI departmental agent actually does	65
8.6 What Anicca runs today	66
8.7 What is an AI operating system?	67
8.8 Examples of existing AI operating system products	68
8.9 Claude for Small Business	69
8.10 The technology infrastructure behind company AI	70

CHAPTER 8 (CONTINUED)

8.11 Build it yourself, buy a system, or use a guided approach?	71
8.12 The economics of a company AI operating system	72
8.13 SME implementation roadblocks	73
8.14 Why this becomes change management, not just AI training	74
8.15 The C10 framework as the diagnostic	75
8.16 Introducing C ⁴ : Company Command & Control Centre	76
8.17 How C ⁴ came together	77
8.18 What C ⁴ looks like	78
8.19 What is different about C ⁴	79
8.20 Where to start	80

APPENDICES

Glossary (A-G)	81
Glossary (H-Z)	82
Appendix A: Setting Up Your Permissions	83
Appendix B: What to Say to Claude	85
Your Next Step: The Thursday AI Club	87
Appendix C: References and Further Reading	88

CHAPTER ONE

01

Why Claude Code Matters Now

Why the leaders who act now will be the ones setting the pace for the next decade

EXECUTIVE BRIEFING

88% of companies are using AI somewhere, but barely 10% have scaled it to deliver real value. The gap is not caused by bad technology. It is caused by leaders who have not yet seen what AI can actually do beyond the chatbot on their homepage.

That is about to change - and the businesses that move now will be the ones defining the new standard for everyone else.

This chapter gives you the context, the language, and the evidence to understand why this moment is different from every AI announcement that came before it - and why the decisions you make in the next twelve months will matter more than most.

16%

Of UK businesses
currently using AI
DSIT, 2025

80%

Of business AI use
is still basic chat
OpenAI, 2025

75%

Report improved
productivity with AI
DSIT, 2025

1.1 What Is Claude Code and Why Agentic AI Is Changing How Businesses Operate

I have spent nearly three months working in Claude Code every day, and it has completely changed the way I work. I have built a library of reusable AI skills covering competitor intelligence, content production, client reporting, branded ebooks, and data dashboards – all without a development team.

What Is Claude Code?

An **AI-powered operating environment** that runs on your computer. Not a chatbot. Not a single agent. A platform combining all seven capabilities below into one place.

You talk to it in plain English. It reads your files, remembers your preferences, follows your processes, and produces finished work directly on your machine.

What Is an AI Agent?

Software that can **plan, decide, and act** on your behalf, either semi-autonomously (with your review at key steps) or fully autonomously on a schedule.

Unlike a chatbot that waits for each question, an agent reads your brief, breaks it into steps, executes them, and adapts when things change. See sections 1.4 and 5.1.

To explain what makes this possible, there are seven core elements you need to understand:

1 Conversational Interface

Instead of typing instructions, you can talk to it like an assistant using built-in voice, system dictation or your preferred voice-to-text tool. Paste links, upload images, share files. No coding needed – just a conversation.

2 Context

It reads your actual files and folders – brand guidelines, templates, previous work – not just what you paste into a chat window.

3 Memory

Remembers decisions, preferences, and project context between sessions, so it works the way your business works.

4 Skills

Reusable instructions that capture exactly how your organisation works. Build once, use forever across your whole team.

5 Agents

AI that takes action on your behalf – writing reports, publishing content, analysing data – rather than simply answering questions.

6 Tools & APIs

Reads, writes, and manipulates your files directly. Connect it to your CRM, email, databases, and third-party software.

7 Outputs

Not just code. Finished documents, reports, dashboards, websites, branded PDFs, slide decks, and images – ready to send to a client or present to a board.

1.2 The Adoption Gap

88%

of companies use AI
but only 10% at scale

Adobe: AI Inflection Point, 2025

80%

of business AI use
is still basic chat

OpenAI Enterprise Data, 2025

16%

of UK businesses
currently using AI

DSIT UK Business AI Survey, 2025

75%

report improved
productivity with AI

DSIT UK Business AI Survey, 2025

That gap is not caused by bad technology. It is caused by a lack of clarity at the top about what agentic AI actually is and how it differs from the chatbot everyone already has. The problem is not whether AI works - it is that most organisations have not moved beyond the basics.

The real gains do not come from asking ChatGPT better questions. They come from building **AI agents** that read your files, follow your processes, and produce finished outputs without you sitting at the keyboard.


"I spoke to over 4,000 businesses in 2025. Around 80% were just using ChatGPT. Most had not heard of Claude. When I mentioned AI agents, fewer than 15% had a clue what I was talking about. Adoption is slow - not because the tech is not ready, but because it requires the will and the time to learn."

Ann Stanley - Founder & CTO, Anicca Digital

This guide exists to close that gap. The next page shows you where most businesses sit today - and where the real value starts.

1.3 Where Are You on the AI Adoption Ladder?

Most businesses are stuck on Rung 2 - using ChatGPT for basic prompts and copying outputs into documents. The real productivity gains start at Rung 4 (automated workflows connecting AI to your tools) and Rung 5 (fully interactive AI agents that plan, adapt, and deliver). This guide is about getting you to the top.



1.4 Automated Workflows vs AI Agents: Two Different Approaches

There are two fundamentally different ways to use AI in your business. Understanding the difference will save you months of going down the wrong path.

AUTOMATED WORKFLOWS

n8n, Make, Zapier

You build a workflow by connecting steps visually. It can include AI that makes decisions, and the output can be completely different each time based on the input - even powering voice agents. But the workflow itself is fixed: the steps, the connections, and the logic are all determined by what you built. Changing the process means rebuilding the workflow.

INTERACTIVE AI AGENTS

Claude Code/Cowork, Google Antigravity, Perplexity Computer

You describe what you want in plain English. The AI reads your files, plans the steps, makes judgement calls, and adapts. Like briefing a team member - you explain the outcome, they figure out how to deliver it, and they improve every time.

This guide focuses on Claude Code (see section 1.6 for details of all three platforms).

CONSIDERATION	AUTOMATED WORKFLOW (N8N)	AI AGENT (CLAUDE CODE)
How you build it	Drag and connect nodes on a visual canvas	Describe what you want in plain English
What it follows	A set map of routes - like the underground, it follows the tracks but takes different paths depending on the inputs	Its own judgement, adapting to each situation
When it breaks	You debug the connectors and re-wire	You tell it what went wrong, it fixes itself
Best for	Simple, repetitive triggers (form to email, data sync)	Knowledge work (reports, proposals, analysis, content)
Who can build it	Someone comfortable with technical tools	Anyone who can describe what they want
Time to first result	Weeks of learning	Hours or days of learning

BOTH HAVE THEIR PLACE

At Anicca, we use n8n for scheduled data pulls and simple triggers. But for knowledge work - reports, analysis, proposals, content - interactive AI agents are faster and accessible to the whole team.

1.5 How Humans and AI Work Together

Not every task needs the same level of human involvement. The best results come from matching the right operating mode to the right task.

Match the mode to the task. Quick lookups and simple questions are Mode 1 (ask and use). The structured brief-draft-review cycle is Mode 2 (the Human-AI Sandwich). Complex creative work like reports and audits thrives in Mode 3 (the Human-AI Salad), where you and AI collaborate iteratively. Price monitoring and email triage run unsupervised in Mode 4. The skill is knowing which mode fits which task.

MODE 1

Human Accepts AI Output

You ask a question or give a simple instruction. The AI responds and you use the answer. Quick, single-shot, no iteration.

MODE 2

The Human-AI Sandwich

Human sets the brief, AI produces the draft, human reviews and refines. A structured flow: Human → AI → Human.

MODE 3

The Human-AI Salad

Human and AI work interactively, blended together. You contribute ideas, context and direction throughout. Iterative and conversational, not sequential.


MODE 4

Full Automation

The AI runs on a schedule with no human involvement. It monitors, processes, and acts. You step in only if something goes wrong.

Examples of all four modes in practice are shown in later chapters.

The image below illustrates the two collaborative modes - Sandwich (Mode 2) and Salad (Mode 3) - which are where most day-to-day work happens.



1.6 Why Claude Code?

Claude Code is not the only AI-powered operating environment. Google launched **Antigravity** (a browser-based code editor), and Perplexity released **Computer** (a multi-model orchestrator that coordinates across 19 AI models). All three aim to let non-developers build with AI, but they take very different approaches.

CONCEPT	CLAUDE CODE	GOOGLE ANTIGRAVITY	PERPLEXITY COMPUTER
Approach & Access	Desktop app, VS Code/Cursor extension, or web app. Files stay on your machine	Browser-based code editor (IDE). No install needed	Cloud app across 19 models. No local files
Context	CLAUDE.md - editable file you control	Rules files - similar concept, editable	Server-side preferences - not editable
Skills	SKILL.md files, slash commands	SKILL.md files - very similar to Claude	Saved workflows, created conversationally
Memory	MEMORY.md - transparent, editable	Knowledge Items - persists, less visible	Server-side, no user-editable file
Agents	Full local automation, scheduled tasks	Multi-agent orchestration in browser	Strong - orchestrates across apps
Outputs	Files on your machine - full control	Files in browser workspace	Cloud sandbox - no local file access
Models	Anthropic only (Claude Opus, Sonnet)	Multi-model (Gemini, Claude, OpenAI)	19 models orchestrated automatically
Cost	\$20-200/month (Pro \$20, Max \$100 or \$200)	Free preview, then \$20-250/month	\$200/month (Perplexity Max)

WHY THIS GUIDE FOCUSES ON CLAUDE CODE

We have built over a hundred reusable skills and delivered real client work with Claude Code every day. Everything in this guide comes from that experience.

Your files stay on your machine - nothing is locked in a cloud service. You teach it how your business works using plain files you can see and edit, not hidden settings. It works however you prefer: desktop app, Cursor or VS Code, web browser, or command-line terminal (used by developers).

The result is a repeatable process: same input, same quality output, every time. The architecture covered here is becoming the industry standard - what you learn transfers directly to other platforms.

By April 2026, according to Pragmatic Engineer's February 2026 survey of around 15,000 developers, 71% of those regularly using AI agents chose Claude Code, and three out of four developers at startups and smaller companies adopted it as their primary tool. Most teams use Claude Code alongside a code editor (like Cursor or VS Code) for viewing files, and GitHub for backing up everything safely.

"Cursor is the wrapper. Claude Code is the brain. GitHub is our Dropbox for AI."

Darren Wynn - Managing Director, Anicca Digital

1.7 Who Should Implement First

The instinct is to start with a pilot team or a junior hire who "has time to experiment." This is almost always wrong. The people who should use Claude Code first are the ones closest to the decisions, the deadlines, and the bottlenecks.

Start with the People Who Feel the Pain

- **Create documents under time pressure** - proposals, reports, audit decks, client deliverables
- **Repeat the same process with variations** - monthly reporting, competitor analysis, onboarding packs
- **Bridge departments** - translating data into narratives, turning strategy into action plans
- **Have domain expertise but limited tools** - they know what good looks like but spend hours building it manually

THE AMBASSADOR MODEL

Identify one or two ambassadors per department. These are people who are curious, not necessarily technical. Give them structured training, a real business problem, and 30 minutes a day for the first two weeks.

Their role is to build the first working skills that others can use. When a sceptic sees a colleague produce a branded 40-page report in 90 minutes, the conversation changes.

Recommended First Movers by Role

ROLE	WHY THEY SHOULD GO FIRST	FIRST WIN
Directors / Founders	Set the vision, see the strategic value, unblock budget	CLAUDE.md that captures business context
Marketing Leads	High volume of repetitive content and reporting	Competitor analysis or campaign report skill
Operations Managers	Process documentation, SOPs, data transformation	Meeting notes to action items skill
Account Managers	Client-facing documents on tight deadlines	Proposal or QBR deck skill
Hands-on Managers	Close to both strategy and execution	Dashboard or data summary skill

The Adoption Curve in Practice

Innovators (2-3%)

Already experimenting. Give them space and ask them to document what works.

Early Adopters (13-15%)

Your ambassadors. Pair with innovators, give them a real problem in week one.

Early Majority (34%)

Need proof. They adopt when they see colleagues producing better work faster.

Late Majority + Laggards

Will come when it becomes normal. Build the infrastructure, let results convince.

CHAPTER TWO

02

Understanding Claude Code

What Claude Code is, where it runs, and what is changing right now

EXECUTIVE BRIEFING

The biggest misconception about Claude Code is in the name. It is not about coding. It is about organising your thinking into repeatable, improvable instructions that an AI can execute.

Start with Claude Projects on the web – structured instructions that produce consistent output, no installation needed. When you want AI working directly on your machine, download Claude Desktop which includes Cowork for real-time collaboration with your files. For advanced workflows – multi-tab editing, hooks, headless automation, the Agent SDK, or simply the preference to work inside a code editor – move to Claude Code in Cursor or VS Code. Some of this functionality is also emerging in Microsoft Copilot.

5

Essential skills
to build first

0

Lines of code
required

1 week

From install to
first deliverable

2.1 The Evolution from Projects to Claude Code

Most people start with **Claude Projects** on the web at claude.ai, and that is exactly the right place to begin. No installation, no code editor, no technical setup. You upload documents, write instructions, and get structured output. It works immediately.

Over time, you will want more. You will want Claude to edit files on your computer, build branded documents, connect to your business tools, and remember how you work. That is when you move to Claude Code. It is not a replacement for Projects - it is the next step.

Even after months of daily Claude Code use, there are tasks where Projects is still the better choice: preparing research from uploaded PDFs, heavy document analysis, or quick questions that do not need your full working environment. The two work together.

I WANT TO...	CLAUDE PROJECTS (WEB)	CLAUDE DESKTOP (CHAT / COWORK / CODE TAB)	CLAUDE CODE IN IDE (CURSOR / VS CODE)	CLAUDE CODE IN TERMINAL
Upload a PDF and ask questions	✓ Yes	✓ Yes	✓ Yes	✓ Yes
Edit files on my computer	✗ No (web only, cannot reach your files)	Chat tab: ✗ no; Cowork + Code tab: ✓ yes	✓ Yes	✓ Yes
Set persistent instructions	✓ Custom Instructions (project-level)	Chat: ✓ styles / profile; Cowork: ✓ global + folder; Code tab: ✓ CLAUDE.md + memory	✓ CLAUDE.md + Claude Code memory	✓ CLAUDE.md + Claude Code memory
Use reusable skills	✓ Yes (Claude Skills)	Chat / Cowork: Claude Skills; Code tab: file-based Claude Code skills	✓ Yes (file-based Claude Code skills)	✓ Yes (file-based Claude Code skills)
Connect to other apps	✓ Pre-built connectors + remote MCP	✓ Remote connectors + local desktop extensions + MCP (Code tab)	✓ MCP servers	✓ MCP servers
Use API keys and OAuth tokens	✗ No local .env or tokens	Chat / Cowork: ✗ connector OAuth only; Code tab: ✓ env vars + .env files	✓ Yes (env vars, .env, OAuth)	✓ Yes (env vars, .env, OAuth)
Automate multi-step tasks	~ Limited (chat only)	Cowork / Code tab: ✓ yes (incl. scheduled)	✓ Yes	✓ Yes (fullest control)
Run scripts and commands	~ Limited sandboxed (no local)	Chat: ~ sandboxed; Cowork / Code tab: ✓ yes (on your machine / VM)	✓ Yes (some CLI features not exposed in IDE)	✓ Yes
Search the web and browse	✓ Yes	✓ Yes	✓ Yes	✓ Yes
Setup required	✓ Sign in on the web	✓ Install Claude Desktop	✓ Install Cursor or VS Code + Claude Code CLI + Claude Code extension	✓ Install Claude Code CLI, run from your Terminal (no code editor needed)

Some of the terms in this table (skills, MCP servers, OAuth, CLAUDE.md) are explained in detail in later chapters.

THE NATURAL PROGRESSION

Start with Claude Projects for quick wins. Move to Claude Desktop when you want AI working directly on your machine, the Code tab and Cowork both run on your local files with skills, memory, and full automation. Graduate to Cursor, VS Code or one of the other advanced routes (terminal, CI/CD, the Agent SDK) only if you need hooks, headless workflows or scripted runs. Most people make this journey in 2-4 weeks.

2.2 Claude Desktop: The Simplest Way In

Anthropic rebuilt the Claude Desktop app in April 2026 with a fundamental structural change: the app now has **three distinct modes**, toggled from three small icons in the top-left corner of the sidebar. Each mode has its own interface, its own navigation menu, and its own default model. You switch between them with a single click.

Chat

The standard Claude conversation most people already know. Write, Learn, and Code shortcuts sit at the bottom of the input area. Runs on the current Sonnet model by default (check the model selector in your Claude app as defaults change). This mode is for quick questions, brainstorming, and general-purpose AI assistance without file access or agentic capabilities.

Cowork

The agentic mode. Cowork has built-in voice input (no separate voice-to-text tool needed), project folder access, scheduling via the Scheduled menu, and Dispatch for delegating work to other agents. Runs on the current Opus model by default (check the model selector in your Claude app as defaults change). Point Claude at any folder and it reads your files, templates, and brand guidelines as context. Toggle Auto Accept at bottom left so Claude applies file changes automatically. This is where your team does real work: building reports, processing data, running audits, and managing multi-step workflows.

Code

The full Claude Code environment, wrapped inside the Desktop app. This gives you sessions, Routines, worktrees, and pinned projects. Runs on the current Opus model by default (check the model selector in your Claude app as defaults change). Power users and developers live here, but most marketing and management staff will rarely need it.

A Note on Naming

Claude Projects (web version at claude.ai) – chat-based, document-friendly, runs in your browser.

Claude Desktop – downloadable app with three modes:

Chat – conversation

Cowork – autonomous task completion

Code – the agentic tool

Claude Code is the same agentic tool found in Desktop's Code mode. You can run it in different setups (or wrappers):

Desktop – the app on your laptop

An IDE (Integrated Development Environment) – a word processor for code, such as Cursor, VS Code, or Google Antigravity

The Terminal – a simple text window that runs commands

It can also run on the cloud (Anthropic's hosted environment), not just on your local machine.


Throughout this guide, "Claude Code" means the agentic tool regardless of the wrapper you access it from.

These setups have different features and benefits – compared in section 2.4.

The practical effect is that Claude Desktop now does the one thing people were still using Cursor for: running multiple parallel AI sessions while you work on something else. Hooks, headless mode, CI/CD and the Agent SDK remain advanced Claude Code workflows, best suited to IDE, terminal or CI/CD setups rather than the beginner Desktop route. For everyone else, the reason to install a code editor just to use Claude Code has largely evaporated.

The same core ideas – skills, CLAUDE.md, MCP servers and credentials – carry across the different setups, although the exact configuration differs by interface. Non-technical staff no longer need to install an extension host or set up a light theme. Voice input now comes built in across Cowork, Cursor and the Desktop Code tab, so you do not need a paid voice-to-text subscription either.

2.2 Claude Desktop: Three Modes, One App



CHAT Conversation

The standard Claude chat. Quick questions, brainstorming, writing. Runs on the current Sonnet model.

COWORK Agentic Work

Built-in voice, project folders, scheduling, and Dispatch. Runs on the current Opus model. Where your team does real work.

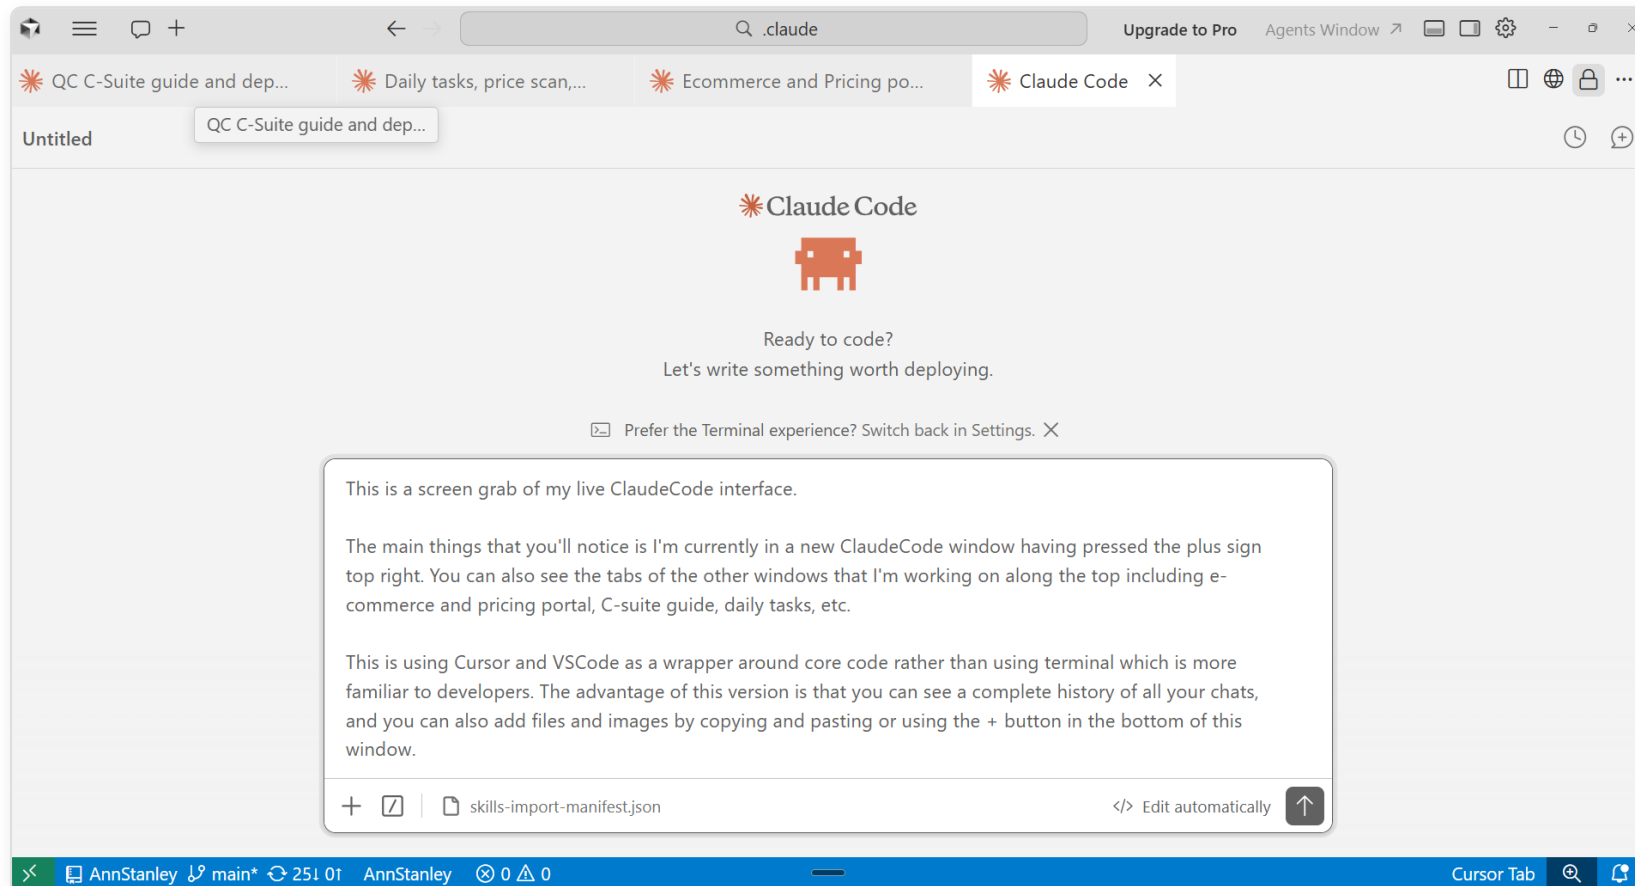
CODE Claude Code Wrapper

The full Claude Code environment. Sessions, Routines, worktrees, pinned projects. Runs on the current Opus model.

The mode toggle sits in the top-left corner of the sidebar (highlighted cyan). Click any icon to switch — your work persists across modes.

2.3 Claude in Cursor / VS Code

A live screenshot of my workspace. Claude Code is running inside **Cursor**, with a white theme, project tabs, and the ability to paste images directly into the conversation.



Tabs across the top - each tab is a separate project. Clock icon pulls back recent sessions.

The conversation - type what you want, Claude acts. Add files and images with the + button.

No code required - the code editor is just the container. Everything here was built in plain English.

A note on jargon: You might hear developers call tools like Cursor an **IDE** (Integrated Development Environment) – think of it as a word processor for code, with built-in features for writing, testing and running software. Other examples include VS Code (Microsoft's free code editor, the industry standard) and Google Antigravity (a browser-based agentic IDE with Gemini built in).

Claude Desktop is *not* an IDE – it is a regular app with a Code tab bolted onto it.

2.4 Choosing Your Setup

Claude Code runs on the same engine regardless of how you access it. The difference is in the interface and a few extras that suit different working styles.

FEATURE	CLAUDE DESKTOP	CLAUDE CODE VIA AN IDE (CURSOR / VS CODE)
Setup	✓ Download app, click Code tab, select folder	✓ Install extension inside your code editor
Best for	✓ Non-technical users, marketers, managers	✓ Teams already working in a code editor
Local files + .env	✓ Read, write, and access credentials	✓ Read, write, and access credentials
MCP servers	✓ Playwright, Chrome DevTools, Gmail	✓ Playwright, Chrome DevTools, Gmail
Git and GitHub	✓ Full version control and collaboration	✓ Full version control and collaboration
Skills + CLAUDE.md	✓ Reusable instructions and memory	✓ Reusable instructions and memory
Deploy	✓ Netlify, Vercel, and other platforms	✓ Netlify, Vercel, and other platforms
Visual diffs	✓ Built-in diff review with inline comments	✓ Full editor diff with syntax highlighting
Live preview	✓ Embedded browser shows your app as you build	✓ Use your own browser
Connectors	✓ Graphical setup for Linear, Notion, Calendar	✓ Configure MCP servers in JSON
Mobile and messaging	✓ Dispatch - send tasks from your phone, review later	✓ Channels - send instructions via Telegram, Discord or iMessage
Voice input	✓ Built-in voice mode in Cowork and the Desktop Code tab	✓ Built-in voice in Cursor. Wispr Flow (\$8-13/month) only if you want the same hotkey everywhere
Parallel sessions	✓ Multiple sessions in a sidebar with side-chat shortcut (April 2026 redesign)	✓ Multiple files + sessions, full file tree, search across project
Hooks	✗ Not available	✓ Automated behaviours triggered by events
Headless mode	✗ Not available	✓ Scriptable for CI/CD and GitHub Actions
Agent SDK	✗ Not available	✓ Build custom agents programmatically

OUR RECOMMENDATION

Start with **Claude Desktop**. It works like a normal app, like Word or Outlook. No code editor, no extensions, no technical setup. The core Claude Code workflows in this guide all work there. Move to **Cursor or VS Code** (an IDE) later if you need **hooks**, **headless mode**, the **Agent SDK**, or prefer working inside a code editor. The April 2026 Desktop redesign added parallel sessions and a side-chat shortcut, so the multiple-chats reason for using Cursor no longer applies. Many teams use both.

Other routes: Claude Code is also available through web, JetBrains IDEs, Slack and CI/CD routes. This guide focuses on Claude Desktop and Cursor / VS Code because they are the most relevant starting points for marketers, managers and business users.

A note on Terminal: Claude Code also runs directly in the Terminal, a plain text window designed for developers. We do not cover it in detail here because it is less suitable for marketers and business owners, who need a visual interface to add examples, preview documents, and review output as they work.

2.5 The Innovation Edge

The Agent Operating System - What Changed in March-May 2026

FEATURE	PLATFORM	MONTH	WHAT IT DOES	WHERE IT HELPS
Claude for Small Business	Web + Desktop	13 May 2026	A small-business package with 15 ready-to-run workflows, 15 skills, and connectors to QuickBooks, PayPal, HubSpot, Canva, DocuSign, Google Workspace and Microsoft 365	Makes AI useful inside tools small businesses already use
Orchestration and Dreaming	Managed Agents	6 May 2026	Orchestration lets one agent co-ordinate other agents. Dreaming lets agents propose their own improvements between runs. Outcomes adds self-reporting on performance. The first signs of agents that improve themselves over time	Teams running multi-step automated workflows that need to keep getting better
Claude Design	Anthropic Labs	17 April 2026	Creates polished designs, prototypes, slides, one-pagers and marketing collateral. Exports to Canva, PDF, PowerPoint and standalone HTML. Powered by Opus 4.7	Marketers and founders turning rough ideas into visual assets
Opus 4.7 + xhigh effort	All	16 April 2026	Stronger coding, agentic workflows, vision and complex professional work. Adds a new xhigh reasoning level for high-stakes tasks. Same pricing as Opus 4.6	Long-running tasks, app building, complex analysis
Desktop Redesign	Desktop	14 April 2026	Parallel sessions in a sidebar, side-chat shortcut (Ctrl+;), drag-and-drop panes for terminal, preview, diff and chat	Multi-track work, replaces the main reason for using Cursor
Routines	Cloud (Anthropic servers)	14 April 2026	Cloud-native scheduled tasks that run on Anthropic infrastructure, your laptop can be off. Hourly minimum, API and GitHub triggers	Daily reports, monitoring, automation
Cowork (GA)	Desktop	9 April 2026	Cowork moved to general availability on Pro, Max, Team and Enterprise. Point it at a folder, give it tasks, get finished work	Knowledge work, documents
Managed Agents (public beta)	Cloud (Anthropic servers)	8 April 2026	Cloud-hosted autonomous agents with sandboxed execution and persistent checkpointing. \$0.08/hr. Notion, Rakuten and Asana were early adopters	Background tasks, CI/CD
Auto Mode	Claude Code	24 March 2026	An AI classifier decides which actions are safe to take without asking you first. Fewer interruptions, faster work. Team plan first	Speed, less interruption
Computer Use (preview)	Desktop + Code	23 March 2026	Claude sees, clicks, and navigates your desktop directly. Research preview, included in Pro and Max at no extra cost. Mac-first	Forms, spreadsheets, apps
Channels	Claude Code	20 March 2026	Send messages via Telegram, Discord, or iMessage. Two-way communication through MCP	Remote control, mobile access
Dispatch	Desktop	17 March 2026	Give Claude a task and walk away. It creates a branch and sends a pull request when done. Delivers results back to your phone	Async work, delegation
Opus 4.6 + 1M Context (GA)	All	13 March 2026	Million-token context window promoted from beta to general availability on Opus 4.6 and Sonnet 4.6. Now superseded by Opus 4.7 (see row above)	Complex analysis
On the Roadmap	Planned	Next	KAIROS (always-on assistant), Coordinator Mode (full multi-agent orchestration), Agent Teams (parallel workers running together)	Future capabilities

We are moving from AI that answers questions to AI that runs processes. The organisations that build their agent operating system now will have a structural advantage that compounds every quarter. Read the latest: anicca.co.uk/blog

2.6 Your First Practical Win

You have seen what Claude Code can do. You understand the different ways to access it. Now let us get you a concrete result in the next 30 minutes.

The Meeting Notes Exercise

This exercise works whether you are using Claude Desktop, VS Code, Cursor, or even Claude Projects on the web. Skills are managed differently in each (account-level Claude Skills in the web and Desktop versions; file-based Claude Code skills in Cursor, VS Code, and Desktop's Code tab), but the exercise itself runs everywhere. It requires no setup beyond having Claude running.

Step by Step

1. **Find your most recent meeting notes** - a Teams transcript, a voice memo, even handwritten notes you can photograph
2. **Paste or attach them to Claude** and say: "Turn these meeting notes into a structured summary with decisions made, action items with owners, and open questions"
3. **Review the output** - Claude will organise the chaos into clear sections. Check it caught the key points
4. **Refine it** - say "Add a section for follow-up dates" or "Make the action items more specific with deadlines"
5. **Save it** - you now have a professional meeting summary that would have taken 20 minutes to write manually

Why This Exercise Matters

Meeting notes are the perfect first task because:

- Everyone has them, so the material is already on your desk
- The quality improvement is immediately visible
- There is no risk - you are not changing any system or process
- It demonstrates the core skill: giving Claude context and getting structured output
- The result is something you can share with a colleague in two minutes

Then Save It as a Skill

Once your meeting summary looks right, tell Claude: "Save this as a skill called /meeting-notes." Next meeting, just type /meeting-notes and paste the transcript. If the output is not quite right, say "improve the skill" and Claude will refine it. By the third use, it will be perfect - and you will notice Claude has created a folder and organised your meeting notes into a sensible record. This is the core loop covered in detail in Chapter 5: do it, save it as a skill, test it, improve it.

"The first time you paste a 40-minute meeting transcript and get a perfect summary in 15 seconds, something clicks. You stop thinking about AI and start thinking about what else you can hand over."

2.6 Your First Practical Win (continued)

Five More Quick Wins for Your First Week

Once you have the meeting notes working, try one of these. Each takes under 10 minutes and requires nothing beyond Claude and a file to work with.

1. Email Draft

Paste a brief and say "Draft a professional email to [recipient] about [topic]. Keep it under 200 words." Edit to taste, send.

2. Data Summary

Attach a spreadsheet and say "Summarise the key trends in this data. What are the three most important takeaways?"

3. Process Documentation

Describe a process you do regularly and say "Turn this into a step-by-step SOP that a new team member could follow."

4. Competitor Quick Look

Give Claude a competitor's URL and say "Visit this website and tell me their main value propositions, target audience, and pricing model."

5. Rewrite for Tone

Paste any piece of writing and say "Rewrite this in a more [professional/casual/persuasive] tone while keeping the same key points." Compare the two versions.

BUILDING YOUR SKILLS LIBRARY

Once you are happy with the output from any of these exercises, follow the same process as the meeting notes: save it as a skill, use it next time, refine it, and repeat. Each one becomes a reusable tool that produces consistent, high-quality output every time. Before long, you will have a library of skills covering the tasks you do most often. Chapter 5 explains the full skills framework, but you can start building yours right now.

CHAPTER THREE

03

Getting Set Up

From first install to a working environment, one afternoon is all it takes

EXECUTIVE BRIEFING

Setting up Claude Code takes one afternoon, not one quarter. This chapter walks through the practical steps: installing the tools, connecting your accounts, understanding the safety controls, and knowing what it costs.

Every setting described here can be changed later. The goal is not perfection on day one, it is getting to a working setup quickly so your team can start building real skills and seeing real results.

1

Afternoon to
full setup

\$20

Starting cost
per month

4

Steps to your
first result

3.1 Setting Up: Day 1

You do not need to be technical. Claude Desktop works like a normal app, like Word or Outlook. Download it, sign in, point it at a folder, and start typing. Our Managing Director Darren set this up himself and recorded tutorials for the team. Here is exactly what he did.

STEP 1: DOWNLOAD CLAUDE DESKTOP (FREE APP)

Go to claude.ai/download and download for your system (Windows or Mac). Install with default settings. If you do not already have an Anthropic account, sign up at claude.ai first. You will need a paid plan to unlock the Code tab – see section 3.4 for costs.

STEP 2: OPEN THE CODE TAB

Claude Desktop has three tabs along the top: **Chat**, **Cowork**, and **Code**. Click **Code** to switch from regular chat mode into the agentic tool. This is the same Claude Code engine developers run from a terminal – you just get a cleaner interface.

STEP 3: POINT CLAUDE AT A FOLDER

Click **Select Folder** and choose a folder on your computer for your first project. Claude will read the files in that folder as context and can create, edit, and organise new files there. Start with a simple folder like *Documents/Claude Projects/Test*. You can change folder any time or open a new session for a different project.

STEP 4: START YOUR FIRST CONVERSATION

Type your first instruction in plain English. For example: *"Look at the files in this folder and tell me what you can see."* Claude will read the folder, describe the contents, and wait for your next instruction. To start a new conversation at any time, open the **sidebar** and click **+ New Session**. Each session is its own project context, and the April 2026 redesign lets you run multiple sessions in parallel.

ALREADY USING CURSOR OR VS CODE?

If you prefer to work inside a code editor, you can install Claude Code as an extension in **Cursor** or **VS Code** instead. **Section 3.5** at the end of this chapter has the four-step setup guide, and section 2.3 shows what the Cursor interface looks like. Graduate to Cursor when you need **hooks** (automation triggered by events), **headless mode** (scripted runs for scheduled tasks), or the **Agent SDK** (building custom agents). For most marketing and management work, Claude Desktop is simpler and has everything you need.

3.2 Additional Setup Including Version Control and Voice Commands

There are additional tools we recommend to complete your setup and improve your experience with Claude Code. GitHub and Git handle version control and backup, Netlify lets you share reports and dashboards on the web, and voice input lets you dictate instead of typing. Claude Desktop Cowork, Cursor and the Desktop Code tab all have built-in voice commands. You only need a separate voice-to-text tool if you are working in Terminal, or if you are already used to a tool like Wispr Flow and want to use it everywhere.

STEP 5: INSTALL A VOICE-TO-TEXT TOOL (OPTIONAL - MAINLY FOR TERMINAL USERS)

Skip this step if you are using Claude Desktop Cowork, Cursor or the Desktop Code tab - all three now have voice commands built in at no extra cost. You only need a separate voice-to-text tool if you are working in Terminal, or if you have already standardised on a tool like **Wispr Flow** (\$8-13/month) and want to use the same hotkey everywhere. Built-in macOS or Windows dictation is free and works well as an alternative. Voice input is roughly three times faster than typing and you naturally provide more context when speaking.

STEP 6: SET UP GITHUB AND GIT

Sign up at github.com (free). GitHub is cloud storage for your projects, organised into **repositories** (repos), think of each repo as a separate project folder in the cloud. Then download **Git** from git-scm.com and install with default settings. Git is the software that runs on your computer and sends your files to GitHub. You will rarely interact with Git directly, Claude Code handles it for you.

STEP 7: SIGN UP FOR NETLIFY

Go to netlify.com and sign up with your GitHub account (free tier is fine). Netlify lets you deploy HTML reports, audit decks, and client portals to the web with a single command. Password protection is built in.

The Importance of GitHub and Git for Version Control and Backup

WITHOUT GIT + GITHUB	WITH GIT + GITHUB
✗ Files only exist on your laptop	✓ Everything backed up to the cloud automatically
✗ No way to undo a bad edit	✓ Roll back to any previous version in seconds
✗ Skills stay on one machine	✓ Share skills across your team

NO TECHNICAL COMMANDS NEEDED

Git is normally controlled through a terminal, the text-based window developers use to give instructions to a computer. You will not need to do this. Claude Code handles all the Git commands for you. Just say *"commit and push to GitHub"* at the end of a session and your work is saved to the cloud.

3.3 Safety First: How Claude Asks Before It Acts

The number one concern from business leaders: **"What if it changes something it should not?"** Claude Code has operating modes that control how much autonomy it has. You can toggle between Plan, Ask Before Edit, and Edit Automatically using the mode button at the bottom right of the prompt box.

The Four Modes for Using Claude Code

PLAN MODE

Think First

Read-only. Claude researches and recommends but cannot change anything on your machine.

ASK BEFORE EDIT

You Approve Each Change

The default. Claude asks permission before every file edit. You click "Yes" each time. Safe but can get tedious.

EDIT AUTOMATICALLY

Files Free, Commands Ask

Claude edits files without asking, but still asks before running commands. A good middle ground once you trust it.

AUTO MODE

Trust and Verify

AI classifier approves routine actions, blocks risky ones. Press Shift+Tab in any Setup to toggle for the session, or set persistent rules in a settings file.

Even with **Edit Automatically** selected, Claude will still ask your permission before running commands such as `bash` (the command line on your computer) or `fetch` (retrieving data from a website or API). This means you will be clicking "Yes" dozens of times per session, which quickly becomes frustrating. **Auto Mode** solves this. Introduced in March 2026, it uses an AI classifier to approve routine actions automatically and only interrupts you for genuinely risky operations. You need to configure it in your settings (or ask Claude to do it for you), but once enabled, it transforms the experience.

Your Data Stays on Your Machine - But Back Up Your Work

In local Desktop, IDE and Terminal workflows, Claude works with files on your own machine rather than locking them inside a separate workspace. You should still follow your organisation's data security rules, especially when working with client, personal or commercially sensitive data. Claude can also edit and delete files if you instruct it to, and it does not have an undo button. Regular backups using GitHub (section 3.2) are essential - if Claude makes a mistake or an irreversible change, you can always roll back to a previous version.

The Dangers of OpenCLAW

You may have heard of OpenCLAW or be tempted to try it. It allows Claude Code to run autonomously without the checks and balances described above, which is attractive to developers because it can run overnight or be controlled via a messaging app. However, we would not recommend it until you have considerable experience. Without proper safeguards, it can make serious mistakes such as deleting all your files. It is also vulnerable to techniques like **prompt injection**, where malicious content in a file tricks the AI into taking unintended actions, which could lead to significant security and compliance issues. To understand what OpenCLAW is, read [The Marketer's Guide to OpenCLAW](#) on the Anicca blog. For enterprise permission controls, see Appendix A.

3.4 Software Tools and Costs

Chapter 3 uses **Claude Desktop** as the default setup. It is the simplest way to get productive with Claude Code and works like a normal app - no extensions, no code editor, no technical setup. Power users who need **hooks**, **headless mode**, or the **Agent SDK** can run Claude Code inside **Cursor** or **VS Code** instead - section 2.3 shows what that looks like. The engine is the same, only the wrapper changes. For the Desktop route you need a Claude subscription, the app, and a folder to point Claude at - that is all. Voice input is built into Desktop Cowork, Cursor and the Desktop Code tab, so most users will not need a separate voice tool. Terminal users, or anyone already used to **Wispr Flow** and wanting one hotkey across every app, can add it for \$8-13/month. **GitHub** and **Git** provide version control and backup, so every change is tracked and recoverable. **Netlify** lets you publish reports, dashboards, and client portals to the web with password protection. **Supabase** is a database tool that only becomes relevant if you are building interactive applications or portals.

ITEM	COST	NOTES
Claude Max subscription	\$100-200/month	claude.ai - Pro (\$20) works but you will hit limits quickly
Cursor or VS Code	Free	cursor.com or code.visualstudio.com
Claude Code extension	Free	Install from VS Code / Cursor marketplace
GitHub account	Free	Version control and backup - essential for team collaboration
Netlify	Free tier available	Deploy reports, portals, and dashboards to the web
Wispr Flow <i>(optional)</i>	\$8-13/month	Voice-to-text for Terminal, or if you already use it everywhere. Cowork, Cursor and the Desktop Code tab all have voice built in
Supabase (optional)	Free tier available	Database for apps and portals - only needed if building tools

The total cost for a single user is \$100-200 per month for Claude Max. Voice input is built into Desktop Cowork, Cursor and the Desktop Code tab, so most users do not need anything extra. Terminal users (or anyone already standardised on Wispr Flow at \$8-13/month) can add the paid voice tool for a total of \$108-213 per month. The remaining tools are free. For a team of five, add GitHub Team (\$4/user/month).

3.5 Setting Up Cursor or VS Code for Power Users

If you need **hooks** (automation triggered by events in `settings.json`), **headless mode** (scripted Claude Code runs for CI/CD or scheduled tasks that read local files), or the **Agent SDK** (building custom agents programmatically), you need to run Claude Code inside a code editor instead of the Desktop app. Here is the four-step Cursor setup Darren recorded for the Anicca team. VS Code works identically if you prefer Microsoft's free editor. The engine is the same as Desktop, only the host application changes.

STEP 1: DOWNLOAD CURSOR (FREE)

Go to cursor.com and download for your system (Windows or Mac). Install with default settings. The free plan is fine, Cursor is just a wrapper for Claude Code so you do not need its built-in AI. Once installed, switch to **light mode** (File > Preferences > Color Theme > Light) as it is easier to read and review documents.

STEP 2: INSTALL THE CLAUDE CODE EXTENSION

Open Cursor. Press **Ctrl+Shift+X** (Windows) or **Cmd+Shift+X** (Mac) to open Extensions. Search for "**Claude**" and install it. You will see a space invader icon appear in the sidebar.

STEP 3: SIGN IN WITH YOUR SUBSCRIPTION

You need an Anthropic account, sign up at claude.ai if you do not have one. Open Claude Code in the sidebar, click "**claude.ai subscription**", copy the link, paste it into your browser, and click Authorize. You may need to upgrade your plan depending on usage (see section 3.4 for costs).

STEP 4: START YOUR FIRST CONVERSATION

Click the space invader icon in the sidebar to open your first chat. To start a new conversation at any time, click the **+** button at the top right, or press **Ctrl+Shift+P** (Windows) / **Cmd+Shift+P** (Mac) and search for "Claude". Claude Code has access to your local files and will create its own project folder structure automatically.

WATCH OUT

Make sure you see the **orange Claude icon** in the sidebar, not Cursor's native AI. If you accidentally prompt Cursor's free tier, it will hit limits quickly. Claude Code through your Anthropic subscription has much higher limits, and on the Max plan you are unlikely to hit them.

CHAPTER FOUR

04

How to Work with Claude

Talk to your AI like you would talk to a colleague

EXECUTIVE BRIEFING

The most powerful feature of Claude Code is also the most overlooked: you talk to it in plain English. No programming language, no special syntax, no training course required.

This chapter covers how to communicate effectively with Claude - through text, voice, images, and files. The better you communicate, the better the output. Think of it less like giving commands to software and more like briefing a capable colleague who needs context, not code.

0

Special syntax
required

3x

Faster with voice
vs typing


80%

First-draft accuracy
with voice skill

4.1 The Seven Concepts That Matter

There are dozens of features in Claude Code, but for a business owner getting started, only seven actually matter. We have expanded Simon Scrapes' original five core concepts to include memory and automation, which are essential once you move beyond your first few sessions.

This guide is built around these seven concepts. Each one has its own chapter, so you can read them in order or jump straight to the one that matters most right now.



The Seven Concepts That Matter - covered across Chapters 4, 5 and 6 of this guide

Expanded from Simon Scrapes' original five core concepts

You do not need to be a developer. I cannot code. Everything I have built - over a hundred skills, audit reports, ebooks, client portals - I have done by describing what I want and iterating with Claude. Occasionally I have needed Claude to help me create JSON files for certain workflows. More recently, when building web apps using Claude Code, it also creates SQL files which I upload into Supabase (a database platform you will learn about later in this guide, used for creating menus, product databases, and similar structured content).

"Every time I need a job doing, I work in Claude Code all day, ask it to do jobs, and then I only come out of Claude Code if I have got a Word document I need to edit."

Ann Stanley - Anicca Digital (March 2026)

4.2 How to Talk to Claude

Brief Claude like you would brief a colleague: be specific, give context, and show examples of what good looks like.

Three Levels of Prompting

LEVEL	WHAT YOU SAY	WHAT YOU GET
✗ Basic	"Write me a blog post about AI"	Generic content that sounds like every other AI output
~ Better	"Write a 1,500-word blog post about SMEs using AI for content marketing, in our brand voice"	Decent first draft, may miss your angle
✓ Best	"Use our AI news blog writer skill (/ai-news-blog-writer) with this week's articles. Match the formatting of the previous news post we created."	Publication-ready, matches your style exactly

The difference between basic and best is not intelligence, it is context. At the "best" level, Claude draws on your saved skills, brand guidelines, and previous examples.

Five Rules

- 1 Be specific about the output.** "Write a 3-slide carousel" beats "create some social content".
- 2 Give examples, not rules.** "Match the tone of this email" works better than a list of tone instructions.
- 3 Iterate, do not restart.** If the first attempt is 70% right, say what to fix rather than starting over.
- 4 One task per conversation.** Start a new tab for each distinct task.
- 5 Say what you would say to a person.** "That is too formal, I would never say leverage" teaches Claude in real time.

You Do Not Have to Type

Claude Code accepts voice, images, screenshots, files, and URLs.

INPUT	HOW	EXAMPLE USE
Voice	Built-in voice in Cowork, Cursor and the Desktop Code tab	Roughly three times faster than typing, brief Claude hands-free
Screenshots	Paste into chat	"Match this design" or "What is wrong here?"
Files	Drag and drop	CSVs, PDFs, Word docs, spreadsheets
URLs	Paste the link	"Scrape this page" or "Analyse this competitor"

4.3 Your Brand Voice

This section is about writing style, not spoken dictation.

Everyone writes differently. Claude does not know how you write unless you show it. A voice skill captures your sentence length, word choices, formatting habits, and the phrases you would never use.

Without one, AI content creates *more* work – you end up rewriting everything. With one, the first draft is already 80–90% right.

STEP BY STEP: BUILD YOUR VOICE SKILL

1. Collect **5-10 real examples** of content you have written (not aspirational – actual posts, emails, reports)
2. **Annotate each one**: what makes it yours? Short sentences? Named people? Specific data? Direct tone?
3. Write a **banned phrases list** – words and clichés you would never use ("unpacking", "game-changer", "here is the deal")
4. Add **formatting preferences**: emoji use, hashtag style, bullet format, call-to-action patterns
5. Define **different post types** if you write different kinds (event promotion vs thought leadership vs news)

PRO TIP: YOUR VOICE SKILL IS A LIVING DOCUMENT

The more examples you give it, the better it gets. When Claude produces something that does not sound right, tell it *why* – then add that feedback to the skill so it does not happen again. My LinkedIn skill knows I write short sentences, name people and companies, and never say "unpacking" – because those are LinkedIn clichés I would never say in person.

4.4 Controlling Thinking Depth

For permission modes (Plan, Standard, Auto), see section 3.3. This section covers how to control how deeply Claude thinks about each request.

Use the `/effort` command to tell Claude how much reasoning to apply:


COMMAND	WHEN TO USE
<code>/effort low</code>	Quick factual answers, simple edits
<code>/effort high</code>	Complex analysis, report writing, debugging
<code>/effort xhigh</code>	Complex coding, multi-step agentic work, app-building, deep restructuring or high-stakes analysis where quality matters more than speed
<code>/effort max</code>	Architecture decisions, restructuring, strategy

Higher effort means Claude spends more time reasoning before responding. For most day-to-day work, the default setting is fine. Save `/effort max` for complex planning tasks.

`/ultraplan` (new, April 2026) - sends your planning task to a dedicated cloud session with up to 30 minutes of reasoning. Your local Claude Code session stays free while it works. You review the plan in your browser, then choose to execute it or bring it back to your local session. Requires a GitHub repository connected.

4.5 Your CLAUDE.md Brand Bible

The single most important file in your setup. CLAUDE.md sits in your project root and tells Claude who you are, what your brand looks like, and how you like things done. Every conversation reads this file automatically.



WHY THIS MATTERS

Once your CLAUDE.md is set up, **everything Claude produces is on-brand from day one**. Build it once, use it everywhere.

HOW OUR TEAM EXPLAINS IT

"A CLAUDE.md file sits in every project you build. When you reload the project, even with a new AI session and no history, it reads that file first and understands what the project is about." - Darren Wynn, MD, Anicca Digital

4.5 CLAUDE.md in Practice (continued)

Here is what a typical CLAUDE.md looks like. Start with these four sections and add more as you discover what Claude needs to know about your business.

```
EXAMPLE CLAUDE.MD (BRAND BIBLE SECTION)
```

```
# My Company Configuration
```

```
## Who I Am
```

```
[Your name], [role] at [company name].  
A [industry] company based in [location].
```

```
## Branding
```

- **Font:** [Your font] (never Arial or Times)
- **Primary colour:** [hex code] (dark backgrounds)
- **Accent colour:** [hex code] (highlights, CTAs)
- **Logo paths:** ~/Documents/Logos/

```
## Tone of Voice
```

- Accessible and practical, not corporate
- Use plain English, avoid jargon
- Write like you are briefing a smart colleague

```
## Document Defaults
```

- A4 landscape for reports and presentations
- Always use branded templates
- Include company logo on every page

Do It. Then Bottle It.

Your CLAUDE.md will grow over time. Every time Claude gets something wrong - wrong font, wrong tone, wrong format - add a rule. After a month, you will have a comprehensive brand bible that prevents the same mistakes from ever happening again.

4.6 Memory, Handovers & Sessions

AI that forgets everything between conversations is frustrating. Claude Code has its own memory behaviour through CLAUDE.md and remembered preferences. At Anicca, we also use a visible MEMORY.md or handover file so the whole team can see project status and decisions.

How Memory Works

Create a MEMORY.md file in your project. Claude reads it at the start of every session. Update it as decisions are made. Think of it as a living project status document that the whole team can read and edit.

EXAMPLE MEMORY.MD

```
# Project Memory
```

Active Projects

Project	Status	Next Step
Cloudfy audit	Draft complete	Ann reviewing
Blog pipeline	Working	Test carousel

Key Decisions

- Use Outfit font for all documents
- Reports are A4 landscape, not portrait

"The more stuff it remembers, the less stuff I have to remember."

Darren Wynn - Anicca Digital Training Session (March 2026)

Handover Documents

At the end of a long session, ask Claude to *"write a handover for this session"*. It captures what was done, what decisions were made, and what is next. The next session picks up exactly where you left off - even weeks later.

4.7 Session Management

"AI context is like milk - it is best served fresh and condensed."

Claude Code community saying (via Simon Scrapes)

Every conversation with Claude has a **context window** - the total amount of text it can hold in memory at once. Think of it as a whiteboard: everything you say, every file Claude reads, every edit it makes gets written on that whiteboard. The current limit is 1 million tokens (roughly 750,000 words, or about 4-5 hours of intensive work), five times larger than the 200K limit in early 2025. As the whiteboard fills up, Claude starts making subtle mistakes - misreading page structures, forgetting earlier instructions, or introducing bugs it would not make in a fresh session. The larger context window helps, but it does not eliminate this problem. This is why session management matters.

Session Discipline

PRACTICE	WHY IT MATTERS
One deliverable per session	Finish one report, one audit, or one piece of work, then write a handover and start fresh. Do not try to do multiple tasks in the same conversation.
Timestamp every message	Add a hook that prints the time at the start of each response. After 90 minutes, you know it is time to wrap up and start fresh.
Name your tabs	Begin your first message by stating the project or task (e.g. "GEO audit for Cloudfy") so Claude auto-names the tab. You can also manually rename tabs in Cursor, VS Code and Desktop code projects (this ability was removed for a while and has come back). A useful pattern is to prefix the day so you can see the order at a glance: <i>MON-Claude guide, TUE-CRE audit, WED-blog</i> . Keep to seven tabs maximum and close old ones after writing a handover.
Run sessions in parallel	Open 3-4 sessions on different tasks. While one builds a report, brief the next on an email. You are never waiting.
Watch the context bar	When it hits 80%, start wrapping up. Do not push to 100% - quality drops sharply in the last 20%.
Big document edits: shorter sessions	HTML reports and ebooks are the most error-prone. After 60 minutes of structural editing, write a handover and start fresh.

WRAPPING UP A SESSION

1. **Write a handover** - say "write a handover" before closing. Claude saves decisions, progress, and what to do next.
2. **Save and commit** - "commit and push" sends your work to GitHub as a safety net.
3. **Start fresh** - new tab, tell Claude which project and to read the handover, then close the old session.

PRO TIP: TALK, DO NOT TYPE

Dictate your messages instead of typing. Claude Desktop Cowork, Cursor and the Desktop Code tab all have voice commands built in. Terminal users (or anyone already used to **Wispr Flow** everywhere) can use a separate voice-to-text tool. You naturally provide more detail when speaking, and better context means better output from Claude.

CHAPTER FIVE

05

Skills: The Practical Operating System

Build once, use forever, share across your team

EXECUTIVE BRIEFING

A skill is a set of instructions that turns a repeatable task into a one-command operation. Your skill library is your competitive moat - it encodes how your business thinks, writes, and makes decisions.

This chapter takes you from understanding what a skill is, through building your first five, to creating systems that improve themselves. The principle throughout: do it manually first, get it right, then bottle it.

100+

Skills built at
Anicca Digital

5

Skills to build
in your first week

7

Levels of skill
maturity

5.1 A Skill Is Just a Folder

Before you build your first skill, you need to understand what one actually is. It is simpler than you think.

your-skill-name/

SKILL.md REQUIRED

scripts/ optional

process_data.py
validate.sh

references/ optional

brand-voice.md
examples/

assets/ optional

report-template.md

THE PROCESS

Step-by-step instructions for a specific task. Written in plain English. This is your SOP - Claude follows it every time.

THE DEEP KNOWLEDGE


Your brand voice. Example outputs you like. API guides and docs. Templates and formats. Executable scripts. Claude only loads these when it needs them.

Context is what makes the process yours.

Drop this folder into your `.claude/skills/` directory and Claude can see and use it. That is all there is to it. The next page shows you how to build one from a conversation.

5.2 The Seven Levels of Skill Building

Skills start simple and get more powerful over time. Most people stop at level 2 or 3. The real competitive advantage comes from levels 4-7, where skills become personalised, measurable, self-improving systems.



Adapted from Simon Scrapes - "7 Levels of Claude Code Skill Building"

5.3 Building Your First Skill: Do It. Then Bottle It.

A skill is not something you sit down and write from scratch. **It emerges from doing real work.** I call this approach *"Do It. Then Bottle It."* - you do the work first, get something you are really happy with, and then bottle it as a reusable skill with all the supporting documents.

DO IT. THEN BOTTLE IT. - THE FOUR STEPS

1. **Do the work** - Use Claude to produce something real: a report, proposal, or blog post. Iterate conversationally until the output is exactly right
2. **Iterate until you are happy** - Refine formatting, tone, and structure. Compare it against existing documents your team has produced. Get it client-ready
3. **Save it as a skill** - Say *"Save this as a skill."* Claude uses its built-in skill-creator to capture the process, brand rules, and output format into a reusable skill folder
4. **Use it again** - Type `/my-skill-name` and Claude follows those same instructions. Same quality, every time

ALTERNATIVE: FIND AND COMBINE EXISTING SKILLS WITH YOUR OWN WORK

You do not have to build everything from scratch. There are skills already built for auditing processes, document formatting, departmental workflows, and much more. Ask Claude to search for relevant skills, or look yourself:

- **Claude's skills database** and GitHub repositories shared by the community
- **GitHub**, where developers share open-source skills
- **LinkedIn**, where people share skills in posts, articles, and newsletters
- **Online directories** and curated skill libraries that are starting to appear
- **Training communities** such as Skool groups run by thought leaders in AI, where members share skills as part of their courses

Once you have found something relevant, combine it with your own work. If you have existing documents your team has produced (a Word audit, a proposal template, an email format), show them to Claude and say: *"Take the best of this existing skill and improve what I have already got."* Work with Claude iteratively until you are happy, then say: *"Build me a new skill that combines the best of both."*

For example, our Tech SEO Audit skill takes James Allen's 50-page technical audit (originally a Word document) and applies our presentation skill to it, producing a professionally designed report that is far better than either input on its own.

PRO TIP: MULTI-TASK WHILE CLAUDE WORKS

Each conversation takes time. Run **5-7 tabs at once**, each on a different task. While one tab builds a report, start another. The time Claude takes on each step is time you can spend reviewing or briefing something else.

5.4 The Skill Creator: The Skill That Builds All Others

This is the meta-skill. It teaches Claude how to write new skills. Once you have it, you can say *"create a skill for writing case studies"* and Claude produces a properly structured SKILL.md you can use immediately.

Progressive Disclosure: How Skills Load Intelligently

Skills use a three-tier loading system. Claude does not read everything at once - it loads information only when needed, keeping the context window lean:

Tier 1

Always Loaded

Every skill file starts with a small header block (called frontmatter) containing the skill name and a short description. This is what Claude reads first to decide whether to use a particular skill. Always loaded, max 15,000 characters.

Tier 2

Loaded When Triggered

The SKILL.md body - your step-by-step process. Keep under 200 lines. This is your SOP.

Tier 3

Loaded on Demand

Reference files, scripts, assets. Claude only pulls these when a specific step requires them.

This three-tier system means Claude is not loading everything at once. It is being smart about what information it brings in and when - which directly improves output quality.

HOW TO: CONVERT EXISTING WORK INTO SKILLS

You do not have to start from scratch. Convert what you already have:

- **Custom GPTs** - Copy your GPT instructions and ask Claude to convert them into SKILL.md format. Most translate directly
- **Claude Projects** - Package your custom instructions and knowledge files as a downloadable skill
- **Productive chat histories** - Had a conversation where you refined a process? Get it summarised and converted into a skill. All that learning is sitting in your chat history

This is the fastest way to build your skill library. Package work you have already done.

5.5 Reproducible Intelligence: Why Skills Are Your Competitive Moat

The real power of agentic AI is not the AI itself. It is the instructions you give it.

Every skill you build captures organisational knowledge in a format that produces **consistent, predictable output**. Your brand bible ensures documents look right, your voice skill ensures posts sound right, your audit skill ensures reports follow the same structure.

The Difference

APPROACH	RESULT	CONSISTENCY
Chat with AI	Different output every time	Varies every time
GPT / Custom prompt	Similar output, but drifts over time	Similar but drifts
Skill: structured content and data (blogs, posts, emails, proposals, data extraction)	Consistent quality from predictable inputs	Reliable
Skill: complex visual documents (HTML reports, audits, slide decks with images)	Strong starting point, needs human QC	Needs final review

A chatbot gives different answers every conversation. Skills with structured inputs (blogs, emails, proposals) give consistent results because the process is predictable. Complex visual documents like HTML reports get you most of the way there, but layout and image sizing still need a human eye. The skill handles the complexity - you handle the final 10-20%.

*In this guide we now use the term **AI departmental agent** for the same idea: a named role made of chained skills with a cadence and a supervisor.*

Your Skill Library Is Your Moat

Every skill encodes decisions your team has made: tone, layout, data, what to leave out. These compound over time. A competitor can buy the same AI subscription, but they cannot buy your skills library. This is the shift from chatbot to **AI-powered operating system** - one that gets better with every conversation, every correction, and every new skill.

"By stringing skills together, what you end up with is an AI employee."

Ann Stanley - Anicca Digital Training Session (March 2026)

5.6 The HTML-First Workflow

One thing I have learned from building all of these skills: Claude is exceptionally good at creating HTML. Beautiful, pixel-perfect pages that can be a report, a carousel, a slide deck, or even a marketing image.

My approach now is to always **get the HTML version perfect first** - layout, colours, fonts, spacing, everything exactly right - and then convert to whatever format I need.

HOW TO: THE HTML-FIRST PROCESS

1. Ask Claude to build the output as **HTML** (not Word, not PDF)
2. Open the HTML file in your browser - check layout, fonts, spacing
3. Iterate until it is **pixel-perfect**
4. Then convert: HTML → PDF (print), HTML → Word (Pandoc), HTML → PNG (screenshot)

This works because HTML gives you complete control over every visual element. Once it is right, the conversion is mechanical.

Organising Your Files

Claude Code reads from your computer, not from uploads. You need a sensible folder structure so Claude can find what it needs.

- Keep client work in **named folders** (e.g., CLIENT 2026/Cloudfy/)
- Keep logos, brand assets, and templates in **known locations** referenced in your CLAUDE.md
- **Right-click any file** → **Copy Path**, then paste it into your Claude conversation
- Use **base64 encoding** for images - embedded images ensure colours and proportions stay accurate (no more stretched logos)

Do It. Then Bottle It.

Get the HTML perfect in one conversation. Then save the build script, the CSS, and the template structure as a skill. Next time, you type one command and get the same quality output instantly.

5.7 QC & Critique: Let Claude Check Its Own Work

Every output needs checking. How you QC depends on whether you are reviewing text or visual layout. Claude can help with both, but the approach is different.

1. Text Content

For text (blog posts, emails, proposals), read the output and correct it directly. Copy-paste the exact wording you want: *"Change the second paragraph to this..."* Claude applies it precisely.

You can also ask Claude to check its own work: *"Review this against our brand voice skill and flag anything that does not match."* Claude will catch inconsistencies in tone, terminology, and formatting that you might miss on a quick read.

2. Visual Content

For layout and design (reports, carousels, slide decks), **do not describe what is wrong**. Saying *"move the logo up a bit"* could mean anything. Instead:

1. Keep a **seed image** in your skill folder - the gold standard
2. Say: *"Compare this to the reference and list every difference"*
3. Claude compares pixel by pixel: positions, spacing, fonts, colours
4. You agree or edit its findings, and the next version is much closer

3. SELF-IMPROVING SKILLS

Advanced skills can improve themselves. Add a **learnings file** that captures feedback at the end of each session: what worked well, what you corrected, what to do differently next time. Keep it structured and prune it weekly. Over time, your skills get better with every use - without you manually rewriting the instructions. This is how our blog pipeline went from needing heavy editing to producing publication-ready posts.

USING VOICE FOR QC

Most of your QC feedback will be spoken, not typed. Cowork, Cursor and the Desktop Code tab all have voice built in, so you can talk through what you want changed straight away: *"move the title down, make the image smaller, change the wording in the second paragraph to..."* Your voice prompts will naturally be longer and less precise than typing, and that is fine. You can say *"interpret what I have just said and give me the best version"* or *"reorganise this section and suggest improvements."* Claude will often come up with ideas you had not thought of. This means you work in stages - a rough voice brief, then a couple of rounds of refinement - but the end result is often better than if you had tried to be precise from the start.

5.8 Your First Week Plan

Block out a two-hour slot each day when nobody will interrupt you. As the week progresses and you get more confident, you can work longer sessions, but two focused hours is the best way to start.

DAY 1

Install & Explore

- Follow Steps 1-4 from section 3.1: install Cursor or VS Code, add the Claude Code extension, sign in, and start your first conversation
- Switch to light mode in VS Code (File > Preferences > Color Theme)
- Ask Claude to read a folder of your existing documents

DAY 2

Setup & Brand Bible

- Follow Steps 5-7 from section 3.2: install a voice-to-text tool if you work in Terminal, set up GitHub and Git, sign up for Netlify
- Create your CLAUDE.md brand bible (section 4.5): fonts, colours, logos, tone
- Test it: ask Claude to produce a short branded document and refine

DAY 3

Practice with Real Work

- Choose something you do regularly: a blog post, email, meeting notes, or social post
- Ask Claude to produce it, then practise editing and formatting
- Iterate until it matches your existing examples of what good looks like

DAY 4

Memory & Handovers

- Write your first handover at the end of a session
- Start a new session and tell Claude to read the handover
- Experience the continuity between sessions

DAY 5

Save It as a Skill

- Take what you built on Day 3 and say *"Save this as a skill"*
- Open a fresh session and run the skill to see how reproducible it is
- Iterate on the skill instructions until you are happy with the consistency

THIS IS A MARATHON, NOT A SPRINT

A week gets you started, but expect three to six months to fully rewire your processes. Start with repeatable, high-frequency tasks for quick wins. As those free up your time, move to quality improvements and new capabilities.

5.9 What to Build Next

Once you have got your first five skills working (brand bible, voice, template, memory, pipeline), here is where to go next. Match your most time-consuming regular tasks to the skill type:

IF YOU DO THIS REGULARLY...	BUILD THIS SKILL
Write a weekly blog or newsletter	A content pipeline (discovery → writing → formatting → publishing)
Send proposals or quotes	A proposal generator with standard pricing and terms
Run audits or assessments	A scored audit template with weighted criteria
Create social media content	A carousel or image generator with brand assets
Onboard new clients	A discovery questionnaire that outputs a structured brief
Monitor competitors	A competitor analysis skill that pulls and formats data
Brief your team	A daily briefing skill pulling from calendar, email, and project notes

Do It. Then Bottle It.

If you find yourself typing the same instructions into Claude more than twice, it is time to bottle it as a skill. This principle has driven the creation of over a hundred skills at Anicca Digital. Each one saves time every single time it is used - and the savings compound as the library grows.

PRO TIP: FIND YOUR EARLY ADOPTERS

You do not build a skill library alone. At Anicca, Holly Kelly built the first skill inside Claude Projects. Find your early adopters - they will accelerate everything.

CHAPTER SIX

06

Business Use Cases and Outputs

The range of applications across every department
– and the benefits that matter to the board

EXECUTIVE BRIEFING

The real power of agentic AI is not in any single application. It is in the breadth. Every department has high-volume, medium-complexity tasks where an AI agent can save time, reduce costs, improve quality, or do things that previously required expensive third parties.

200+

AI skills
built
and
customised

8

Departments
covered

£0

Development
team
required

*Source: Anicca Digital internal data,
March 2026 – anicca.co.uk/blog*

6.1 The Benefits Framework

Every agentic AI application delivers value through one or more of these four types. When evaluating where to start, map your opportunities against this framework:

TIME SAVING

Hours Reclaimed

Routine tasks that took hours now take minutes. Blog pipeline: 4–8 hours → under 2 hours. Competitor reports: 1 week → 1 day. Email triage runs continuously in the background. Expect to save 50–80% of the time on repeatable work.

COST SAVING

Budget Released

Work that required agencies or contractors moves in-house. Ebook design: £3–5k → built internally. LinkedIn carousels: £500 each → free from template.

QUALITY IMPROVEMENT

Consistency & Depth

40+ GEO audit reports over two months with no dev team, each one branded and structured to the same standard. A full library of reusable skills means the quality never drifts, even at volume.

PREVIOUSLY IMPOSSIBLE

New Capabilities

60-page competitor intelligence reports, CRO audits with GA4 data, interactive client portals deployed to the web. The depth of research, design, and analysis in each deliverable would previously have required multiple specialists or expensive agencies.

THE LEARNING CURVE IS REAL - BUT REWARDING

Start with **time-saving wins**: repeatable tasks you do every week (email drafts, report formatting, data pulls). These give you fast results and build confidence. Over time, you move to **quality improvements** (branded outputs, deeper analysis) and then **new capabilities** you could not do before. Expect 3–6 months to fully rewire your processes. Pick high-frequency, repeatable tasks first for the quickest return.

HOW TO READ THIS CHAPTER

The following pages show real applications by department. Each example is tagged with the benefit type it delivers. Look for the ones that match your organisation's biggest pain points.

6.2 The Full Department Map

Every department has tasks that agentic AI can handle. The range of pre-built skills and integrations means you are not starting from scratch.

DEPARTMENT	EXAMPLE APPLICATIONS	PRIMARY BENEFIT
Marketing	Blog pipeline, carousels, brand images, GEO audits	Time + Cost
Sales	Proposals, competitor intel, avatar councils, lead enrichment	Time + Quality
Operations	Email triage, meeting notes, order confirmations	Time
Finance	CSV analysis, expense tracking, reconciliation	Quality + Time
HR & Research	Performance coaching, expert councils, market monitoring	Quality
Creative	Blog images, ads, infographics, ebooks	Cost
IT & Development	Web apps, client portals, dashboards, deployments	Previously impossible
Client Services	Portals, audit reports, QC dashboards	Previously impossible

BEYOND THE OFFICE

One of our highest-rated sessions had nothing to do with marketing - comparing 11 mortgage deals and salary sacrifice options for an electric vehicle. If it involves comparing options and making a decision, Claude Code can help.

THE KEY INSIGHT

The highest-value applications are in the **workflows that cross departments** - where a prospect becomes a competitor report becomes a proposal becomes a client portal.

6.3 Marketing & Content

Marketing has the most immediate wins. Every marketing team has repetitive, template-driven work that is perfect for agentic AI.

What We Actually Do at Anicca

- **Weekly blog pipeline** (*Time saving*) - AI discovers trending articles, scrapes content, writes a categorised news blog, uploads to WordPress, generates a LinkedIn post and carousel
- **Branded carousels** (*Cost saving*) - Professional LinkedIn carousel decks from a template in minutes. Previously £500+ each from a designer
- **Blog featured images** (*Cost saving*) - Branded 1200×628 SVG images with headline text, generated on demand
- **GEO audit reports** (*Time saving*) - 40+ reports over two months, previously a time-consuming manual process

REAL NUMBERS - ANICCA DIGITAL, 2026

Anicca's AI in Marketing and Management Weekly Post

BEFORE

1 full day

AFTER

~2 hours

OUTPUT

Blog + image + LinkedIn post + carousel + newsletter

REAL NUMBERS - CHARLES BENTLEY, 2026

Automated Pricing & E-Commerce Dashboard - Daily Data Collection

MERCHANT CENTER

Stock, prices, and competitor benchmarks

COMPETITOR PRICING

Scraped from rival retailers via Apify

REVENUE & PERFORMANCE

Magento, GA4, and Google Ads per product

Fully automated - the client dashboard updates every day with no manual intervention.

Ann Stanley
Founder & CTO

anicca ai & insights

AI in Marketing & Management

7th Apr 2026

Claude subscription
crackdown
Sephora on ChatGPT
Gemma 4 goes
open-source

6.4 Documents, Reports & Knowledge

Every document type your organisation produces regularly can become a skill. Here is what we have built at Anicca - each one produces branded, consistent output every time.

DOCUMENT	WHAT IT DOES	OUTPUT
Order confirmations	Professional confirmation with scope, pricing, and terms	.docx
Proposals	Three types - campaign, discovery, multichannel - with budget tables and timelines	.docx
NDA's	Non-disclosure agreements with correct legal formatting	.docx
LinkedIn carousels	Branded slide decks for news round-ups and weekly updates	HTML → PDF
Blog posts	Full articles with SEO formatting, table of contents, reading time	HTML → WordPress
Audit reports	AI visibility, CRO, competitor, content, technical SEO - slide decks	HTML (landscape)
Ebooks & guides	A4 landscape branded publications - like the one you are reading	HTML → PDF
Blog images	Branded header images with headline text and illustrations	HTML → PNG

CASE STUDY - ANICCA DIGITAL, 2026

The Ebook You are Reading Right Now

TRADITIONAL DESIGN COSTS	TRADITIONAL TIME	WITH CLAUDE CODE
£5,000-10,000	2-3 months	Built in days

A4 landscape, branded typography, print-ready PDF - designed and written using Claude Code. No graphic designer. No InDesign.

6.5 What Types of Content Can Claude Code Produce?

Claude Code does not just write code. It produces finished, branded deliverables ready to send to a client, present to a board, or publish to the web. In many cases it is pulling data from external sources, analysing it, and reformatting it into something useful.

Documents

FORMAT	WHAT IT IS	REAL EXAMPLE
HTML Reports	A4 landscape slide decks with branding, charts, screenshots	GEO audits (40pp), competitor intelligence (61pp)
PDFs	Print-ready documents generated from HTML	Client deliverables, ebooks, storyboards
Word Documents	Branded .docx files with correct fonts and formatting	Proposals, order confirmations, NDAs
PowerPoint	Slide decks for presentations and workshops	Conference talks, client presentations
Ebooks	Multi-chapter publications (HTML, PDF, or EPUB)	This guide, CIO AI Adoption Framework series

Data, Numbers and Analysis

TYPE	WHAT IT DOES	REAL EXAMPLE
API Data Pulls	Fetches data from Google Analytics, Google Ads, Magento	Daily revenue, traffic reports, ad performance
CSV Analysis	Analyses spreadsheets, generates statistics and charts	Google Ads exports, sales data, financial reports
Classification	Categorises large datasets by intent, topic, or criteria	8,962 LinkedIn companies scored by ICP fit
Calculations	ROAS, TACOS, conversion rates, scoring models	Ecommerce profitability, audit scoring (e.g. 37/100)

Images, Content and Video

TYPE	WHAT IT IS	REAL EXAMPLE
Featured Images	Branded blog headers with SVG illustrations	1200x628 images matching brand guidelines
Carousels	Multi-slide visual summaries for LinkedIn	"What a Week" carousel, AI news roundup
Blog Posts	Written, formatted, and uploaded to WordPress as drafts	Weekly AI marketing news, search marketing news
LinkedIn Posts	Written in your voice, newsletters with companion posts	Article promotion, event announcements
Video Scripts	Storyboards, shot lists, and programmatic video creation	Testimonial scripts, explainer videos

6.6 Real Examples: Documents We Have Built

Every deliverable below was built using Claude Code at Anicca Digital. These are not mockups – they are real client work, produced faster and to a higher standard than traditional methods allowed.

DELIVERABLE	FORMAT	WHAT IT CONTAINS
GEO Audit Report	40–page HTML/PDF	AI visibility scores across 6 platforms, knowledge map analysis, content cluster diagrams, buying funnel visualisations, competitive benchmarking. 12 reports built in one week.
Competitor Intelligence	61–page HTML/PDF	Organic and paid analysis, social media SWOT, tech stack comparison, executive summary with strategic recommendations. Data from 15+ sources.
B2C Avatar Council	45–page branded report	AI-generated customer personas with portrait images, competitor scoring matrices, word cloud analysis, Six Thinking Hats evaluation by 19 expert personas.
Google Ads Audit	69–page scored report	85 weighted checks across campaigns, ad groups, keywords, and extensions. Traffic light scoring, GA4 reconciliation, interactive QC review tool.
Client Proposals	Branded .docx	Three proposal types (campaign, discovery, multichannel) with correct Outfit font, Anicca branding, and client-specific pricing.
Weekly News Pipeline	Blog + carousel + image	Automated article discovery, AI-written blog post, branded featured image, LinkedIn carousel – all from a single command chain.
This Guide	70+ page ebook	A4 landscape, embedded images, table of contents, glossary. Four major versions iterated in collaboration with Claude Code.

"The output is not a prompt response. It is a finished document, a deployed website, a branded PDF. That is what makes this different from ChatGPT."

Ann Stanley – Founder & CTO, Anicca Digital


The common thread is repeatability. Once a deliverable has been built and refined, it becomes a skill – a reusable template that produces the same quality every time, for every client.

6.7 Sales & New Business

The goal is not to replace your sales team – it is to eliminate the hours spent on preparation so they can focus on relationships.

- **Lead enrichment and outreach** (*Time saving*) – Company data enriched via Clay and Instantly, with Twain AI generating personalised emails for each prospect
- **Proposal generation** (*Time saving*) – Three proposal types with branding, tables, and budget breakdowns
- **Competitor intelligence** (*Quality*) – 61-page branded reports covering digital presence, content, paid media, and positioning
- **Customer avatar councils** (*Previously impossible*) – AI panel of 20 expert personas deliberating your ideal customer profile
- **GEO audit reports as lead magnets** (*Time saving*) – 40+ AI visibility reports in four weeks, used as personalised outreach with data prospects can also use as clients

Example content cluster from a GEO audit report - E-commerce Marketing Agency



6.8 Research & Strategy


Research and strategy is where agentic AI goes beyond automation into genuine analysis. These tools produce work that would not have been feasible without a team of specialists.

- **Expert councils** (*Previously impossible*) - Virtual panel of six to eight job roles (B2B) or personas (B2C) representing your target audience. We apply Six Thinking Hats and "buyer councils" to review first impressions, website effectiveness, and external factors like reviews and testimonials
- **Market monitoring** (*Time saving*) - Automated newsletter summaries, YouTube transcript analysis, research indexes built and searchable
- **Conversion rate analysis & content audits** (*Quality*) - Full site audit with Google Analytics data, competitive analysis, buyer persona alignment, and MECLABS per-page assessments
- **Competitor intelligence** (*Previously impossible*) - 61-page branded reports covering digital presence, content strategy, paid media, social, and competitive positioning

6.9 Finance & Administration

- **Email triage** (*Time saving*) - AI grades emails 1-5, drafts replies, sends Teams notifications every 30 minutes
- **Meeting transcription** (*Quality*) - Raw recordings → structured documents with action items and owners
- **CSV analysis** (*Previously impossible*) - Drop a spreadsheet, get statistics, correlations, and visualisations


Content by Type — 726 URLs



6.10 Advanced: Web Apps, Portals & Dashboards

This is where agentic AI moves beyond documents into live, interactive applications. You do not need a developer - you describe what you want, iterate with Claude, and deploy to the web. This is sometimes called "vibe coding".

- **Client reporting portal** - 9-tab dashboard with audits, strategy, pricing, e-commerce, and performance data. Password-protected, updated daily via API
- **Training and community portal** - Course content, video library, skill documentation, and learning paths
- **E-commerce dashboards** - Magento revenue, GA4 analytics, and Google Ads data with date filters and trend charts
- **Pricing monitor** - Daily automated scans pulling Merchant Center, competitor prices, and sales data into one view
- **Interactive audit tools** - QC dashboards, approval documents with editable notes, password-gated sharing
- **Full websites** - Entire sites designed, coded, and deployed without a developer



6.11 APIs, MCP and Credentials

This is where Claude Code stops being a writing assistant and starts being a business tool. It can connect to your data sources, take over your browser, or use third-party tools to access websites directly.

CONNECTION METHOD	HOW IT WORKS	WHEN TO USE	EXAMPLES AT ANICCA
Playwright	Competitor screenshots, report PDFs, page audits	Claude launches a headless browser to navigate pages, capture screenshots, and generate PDFs	Screenshots, PDF generation, scraping public pages
Chrome MCP	Waikay scores, Google Ads UI, authenticated portals	Claude connects to your actual Chrome browser, using your logged-in sessions and cookies	Dashboards and tools that need your login
Direct API calls	Magento product feed, GA4 analytics, Waikay SEO scores	Claude writes a script that calls the service using your credentials from .env	When you need specific data or custom logic
MCP servers	Gmail search and drafts, Google Sheets, Google Calendar	Pre-built connector runs in the background. Claude uses it like a native tool	Google services, databases, file systems

What Is an API?

An API (Application Programming Interface) is a structured way for software to talk to other software. When Claude pulls your website traffic from Google Analytics or reads product prices from your ecommerce platform, it is using their API. Every modern business tool has one: Google, Shopify, Magento, HubSpot, Slack, LinkedIn.

To use an API, you need **credentials** - a key or token that proves you have permission. You generate these in each service's settings (usually under "API access" or "Developer"). They look like long random strings of text and are stored in a single .env file on your machine - a plain text file that Claude reads but never uploads or shares.

What Is MCP?

MCP (Model Context Protocol) is an open standard created by Anthropic that gives Claude pre-built connectors to external services. Think of it as a universal adapter: instead of writing custom code for every service, you install an MCP server (a small background program) and Claude gains new capabilities automatically.

The .env File: Your Credential Vault

All credentials live in a single .env file - one file, one location, never committed to GitHub. Each line is a key-value pair: GOOGLE_API_KEY=abc123 . . . When Claude needs access, it reads this file automatically. When you add a new service, you add one line. This is standard practice across the software industry.

WHY APIS AND CREDENTIALS FEEL DIFFICULT FOR NON-TECHNICAL TEAMS

Setting up API credentials is the single biggest barrier for non-technical teams. Each service has its own process - some take 2 minutes, some need OAuth consent screens. Budget half a day the first time. Once done, it rarely needs revisiting.

6.12 From Individual Tasks to Automated Workflows

Individual skills are useful on their own, but the real transformation happens when you chain them into **workflows** – a sequence of tasks that run one after another, producing a finished result with minimal input from you. You set these up once, and they keep running – you just tell Claude what you want to happen and when.

How It Works: My Sunday Content Pipeline

Every Sunday, I prepare the main content piece that goes out on Monday morning. Seven skills run in sequence:

1. **News discovery** finds 30+ trending articles from industry sources
2. **Blog writer** selects the best stories, writes and formats a 2,000-word blog post
3. **Image generator** creates a branded featured image to go with it
4. **WordPress upload** publishes the blog as a draft, ready for review
5. **Carousel builder** creates a 10–12 slide visual summary for social media
6. **LinkedIn post** writes a promotional post in Ann's voice
7. **Newsletter creator** produces a LinkedIn newsletter version with a companion post

I schedule the carousel and LinkedIn post for 8am Monday. Once the main post goes live, I upload the newsletter through LinkedIn's newsletter functionality.

Total hands-on time: about 2 hours. Before this pipeline existed, the same work took two full days.

Fully Automated Tasks That Run Without You

Some tasks run on a schedule without any human involvement at all. To make this work, I adjusted my laptop settings so it operates even when the lid is closed. Here is what runs automatically at Anicca:

- **Email assistant (every 30 minutes)** – monitors Gmail, grades messages using Claude, sends notifications via Microsoft Teams, and drafts calendar-aware replies.
- **Daily backup (6am)** – checks all activities, backs up skills and memory to GitHub, and updates the project records. Everything is saved and version-controlled before the working day starts.
- **Price monitoring and e-commerce dashboard (midday)** – an automated tool connects to the Merchant Center API to collect prices, stock levels, cost data, product IDs, and supplier names. It then sends out a scraping agent to get accurate live prices for any products that are showing as more expensive in Merchant Center. It runs at midday because retailers typically adjust their prices in the morning. It pulls in the previous day's sales from the Magento API, plus Google Analytics and Google Ads data (costs and revenue), which gives us ROAS and TACOS – a more accurate reflection of advertising cost against total site sales. It also collects weather data and sends a daily alert highlighting any significant price changes relative to the previous week. All of this is displayed in a client portal that was also built in Claude Code, bringing together pricing data, e-commerce sales, and performance data from advertising platforms and SEO.

6.13 Building Web Applications with a Database

For simple reports and dashboards, an HTML file is all you need. But when you want to **save data over time, let people log in with their own accounts, or build a proper web application**, you need a database behind it. We use a service called **Supabase** for this.

What Supabase Gives You

Supabase does a similar job to the WordPress backend - it handles user accounts, stores your data, and manages files - but it handles so much more, and Claude Code builds everything for you with no plugins to install or servers to manage.

- 1 Data storage** - Keep structured information in tables: product prices, client records, audit scores, team member details. When you need to make changes, Claude writes an SQL script - you simply copy it into the Supabase SQL editor and hit run. This might be creating a new table for your products, adding new menu items, setting up user permissions, or updating how data is displayed.
- 2 User accounts and logins** - Let different people log in with different levels of access. Your team sees admin features, clients see their own data, and the public sees a landing page.
- 3 File storage** - Upload and serve images, documents, and reports. Like a cloud drive that your application can access and display.
- 4 Live updates** - Data changes appear instantly without refreshing the page. Ideal for monitoring dashboards and pricing tools.

YOU PROBABLY DO NOT NEED THIS YET

Most businesses will never need a database. Start with skills and HTML reports. Only consider Supabase when you find yourself needing saved data, user logins, or a proper application. At Anicca, we use it for the pricing portal, audit insights dashboard, and community portal - but those came months into our journey, not at the beginning.

6.14 Building Real Applications Without a Developer

Claude Code does not just build static web pages. It can build database-driven applications with dashboards, user logins, API integrations, and live deployment, all described in plain English. What makes this possible is a stack of tools that Claude connects together on your behalf.

Data, Menus and Authentication (Supabase)

Every serious application needs somewhere to store data and a way to control who can see it. Supabase provides both. Claude can set up complex table structures, build navigation menus with multiple levels, create authentication systems with user logins and role-based access, and write the logic to manage, save, and retrieve data. You describe what you need. Claude writes the SQL, and you paste it into the Supabase editor. The result is a dynamic, database-driven application with proper tools and controls, not just a web page.

Version Control and Automation (GitHub)

GitHub stores your entire codebase with a complete audit trail, so you can see every change and nothing is ever lost. But it does more than backup. GitHub Actions can run scheduled tasks, automatically fetching data from APIs, generating reports, or triggering rebuilds. Critically, GitHub acts as middleware between your development environment and your live website, so changes flow through a controlled pipeline rather than being pushed directly.

Live Deployment (Netlify and Vercel)

Netlify and Vercel turn your code into a live website. You can buy a custom URL, point it at your deployment, and add password protection or access controls. When you push a change to GitHub, the site rebuilds and redeploys automatically within seconds. This is how secret-agents.ai works: a fully live website, edited through Claude Code or Cursor, synced via GitHub, and deployed instantly through Netlify. Edit the code, and a new version is live within minutes.

How It All Connects

You talk to Claude Code or Cursor. Claude writes the code and pushes it to GitHub. GitHub triggers a deployment to Netlify or Vercel. The live site pulls data from Supabase and external APIs. Scheduled tasks keep dashboards and reports fresh. The entire chain runs without a developer, and every change is tracked. That is the difference between building a website and building a business application.

6.15 Real Examples: Applications We Have Built

These applications run continuously at Anicca Digital, serving the team and clients every day. Each was built entirely through Claude Code.

APPLICATION	TYPE	WHAT IT DOES
Pricing Portal	Live web application	Pulls daily stock levels and prices from the Merchant Center API, then sends scraping agents to find competitors charging less. Combines Magento sales, GA4 traffic, Google Ads spend, and weather data to track performance and spot patterns. Multiple dashboard views with ecommerce analytics.
Client and Reporting Portal	Live web application	Multi-tab portal consolidating all client audit reports, competitor data, SEO metrics, and performance dashboards in one place. Role-based access for team and clients.
Community and Training Portal	Live web application	Content management for Thursday AI Club, with training courses, event scheduling, member access control, and a resource library. Section-based navigation with calendar, sessions, and community views.
Email Assistant	Automated service	Monitors Gmail every 30 minutes, grades messages with Claude, sends Microsoft Teams notifications, drafts calendar-aware replies. 50+ ignore patterns. Runs via Task Scheduler.
Command and Control	Internal tool	175 skills catalogued with search, filtering, and status tracking. File-driven architecture - the central hub for managing all Claude Code capabilities across the agency.

None of these existed three months ago. Every one was built by a marketer talking to an AI in plain English, not by a developer writing code. That is the shift this technology represents.

CHAPTER SEVEN

07

Implementation, Rollout and Governance

Implementation, governance, and getting your team on board

EXECUTIVE BRIEFING

Technology adoption fails for organisational reasons, not technical ones. The companies that extract value from AI invest in change management, set clear governance, and give their teams permission to experiment.

With Claude now available inside Microsoft 365, the adoption barrier has dropped to near zero. The question is not *if* your team will use AI – it is whether you will guide how they use it.

40%

Of agentic AI projects
cancelled due to governance
Gartner, 2026

29%

Of enterprises prepared
to secure AI agents
Gartner, 2026

7.1 Crawl, Walk, Run

This technology moves fast, and the competitive risk is real. If your competitors adopt agentic AI and you do not, the gap will widen quickly. The good news is that setup takes days, not months, and many skills are now available off the shelf. The time investment is in customising them to your business.

CRAWL - MONTH 1

Learn & Experiment

- Complete the Week 1 plan in section 5.8 (Chapter 3 sets up the tools in an afternoon)
- Build your brand bible and first document template
- Leadership personally uses the tools daily

WALK - MONTHS 2-3

Scale & Govern

- Build a skill library, starting with off-the-shelf skills customised to your needs
- Set governance policies and train 2-3 AI champions
- Deliver your first client-ready output built with AI

RUN - MONTH 4+

Automate & Integrate

- Deploy your first autonomous agent in a low-risk area
- Cross-department workflows and background automation
- Building skills is the main investment, but each one compounds

Who Owns What

ROLE	RESPONSIBILITY
CEO	Strategic direction, resource allocation, culture of experimentation
CMO	First mover – marketing has the most immediate use cases
COO	Process mapping, automation candidates, measuring ROI
CTO/CIO	Technical governance, security review, integration architecture
CFO	Business case validation, budget allocation, ROI tracking

Do It. Then Bottle It.

At every stage, the approach is the same: do the work with Claude first, then save it as a reusable skill. Your crawl phase creates your first skills. Your walk phase creates a library. Your run phase creates an AI-powered operating system.

7.2 What Does It Cost Your Organisation?

Section 3.4 covers setup costs for an individual. When you roll this out across a team, some tools move from free tiers to paid plans. Here is the full picture.

ITEM	FIRST USER	EACH ADDITIONAL USER	NOTES
Claude Pro	Good for getting started	\$20/month	\$20/month
Claude Max	Required for serious daily use	\$100-200/month	\$100-200/month
Cursor / VS Code	The IDE Claude Code runs inside	Free	Free
Wispr Flow <i>(optional)</i>	Voice-to-text for Terminal, or if you already use it everywhere. Cowork, Cursor and the Desktop Code tab all have voice built in	\$8-13/month	\$8-13/month
GitHub	Version control, automation, backup	Free	\$4/month (Team plan)
Netlify / Vercel	Multiple live sites, custom domains, access control	\$19/month (Pro)	Shared
Supabase	Database for apps. Required once you have multiple users or apps	\$25/month (Pro)	Shared

The Real Cost: Time and Training

Software is the smaller investment. The real cost is the time your team spends learning, building their first skills, and adapting workflows. Budget for structured training in month one and reduced output while people build competence. This pays back fast - within weeks, not months - but it is an opportunity cost that needs to be planned for, not discovered by accident.

WHICH PLAN DO YOU NEED?

Start everyone on **Pro** (\$20/month). Upgrade to **Max** (\$100-200/month) when someone consistently hits usage limits - that means they are getting value. Do not buy Max for the whole team on day one.

7.3 Your First Month Health Check

After 30 days of using Claude Code, you will have accumulated enough conversation history for Claude to analyse your patterns and tell you what to improve. The `/insights` command does exactly this.

Running `/insights` scans your conversation history and produces an HTML report covering your interaction style, friction points, concrete CLAUDE.md additions, features you are not using, and usage patterns that may be hurting quality.

REAL EXAMPLE: ANICCA DIGITAL

After 51 sessions, `/insights` revealed HTML report edits were causing layout regressions and suggested five CLAUDE.md rules (backup before edits, plan before editing, verify after, never batch changes, forbidden styling patterns). All five were implemented same day.

It also flagged marathon sessions where quality degraded. The fix: 90-minute blocks with handovers. This single change improved output consistency more than any other adjustment.

The Self-Improving Loop

This is the pattern that separates casual users from teams that get compounding returns from AI:

1. **Work normally** for 2-4 weeks, building skills and completing tasks
2. **Run `/insights`** to see what patterns have emerged
3. **Review the suggestions** - some will be obvious, some surprising
4. **Update your CLAUDE.md** with the high-value rules
5. **Repeat** every month or after any major new workflow

Each cycle makes Claude faster and more reliable. Rules from month one prevent errors in month two, freeing up time for harder problems in month three.

WHEN TO RUN /INSIGHTS

After your first month. After onboarding a new team member. After any week where Claude felt slower than usual. The improvements last permanently.

7.4 Ambassadors and Adoption Styles

Innovators Go First

Every successful rollout starts with one or two ambassadors who are curious enough to invest their own time testing and learning. They need permission to experiment, time to fail, and the drive to figure out what works. Their job is to prove the technology delivers results, so everyone else can see it working before they commit.

Different People, Different Styles

PERSON	FOCUS	WHAT THEY HAVE BUILT
Ann (CTO)	Apps, portals, dashboards	over a hundred skills in three months. Pricing portals, audit reports, client dashboards, community platforms
Darren (MD)	Data and reporting	Reporting dashboards aggregating data from multiple sources. Attribution platforms, operational systems
James (Dev)	Always-on AI systems	AI automation workflows, SEO and GEO tools, voice and chat agents deployed on secret-agents.ai

People naturally gravitate to what they already do well. Do not force everyone into the same approach.

Cascading Through the Business

Once ambassadors have built the first working skills, invest in structured training so every team member reaches a basic level of competence. Start each person with tasks that immediately benefit their role, particularly the jobs they find boring or routine. When someone automates a task they do every week, they see the value instantly and want to do more. Skills are shareable, so one person builds it and the whole team uses it. Follow the benefits quadrant: time savers first, then money savers, then do better, then do new.

7.5 Getting Your Team on Board

The biggest barrier to adoption is not the technology. It is getting people to trust it and use it properly.

Address the Objections Honestly

- **"It will take my job"** - AI replaces tasks, not roles. Free people from repetitive work so they focus on judgment and creativity
- **"It hallucinates"** - Yes, but so do humans under time pressure. Build verification into the workflow
- **"I am not technical enough"** - Claude Code works in plain English. If you can write a brief, you can use it
- **"We tried ChatGPT and it was not useful"** - Agentic AI is fundamentally different. Context, memory, and skills mean it understands your business
- **"We do not have time to invest in this"** - The setup takes one afternoon. After that, it saves time every single day. The cost of not starting is higher
- **"The output is not as good as a human"** - Untrained use produces poor results. With proper context, skills, and prompting, the output quality matches or exceeds manual work in a fraction of the time
- **"AI cannot know our business like we do"** - That is exactly what CLAUDE.md, memory, and skills solve. You teach it your brand, your processes, your data sources. It learns and remembers

Quality Depends on Training

Without proper training, people get poor results and lose confidence. If someone asks ChatGPT to "write me a blog" without understanding prompting frameworks like COSTAR or markdown formatting, the output will be generic. Even using pre-built agents like those on secret-agents.ai, where the prompting is structured through a form, produces noticeably better results. The quality gap between a trained and untrained user is enormous.

Invest in Structured Learning

Do not expect people to figure it out alone. Provide structured training covering prompting, tool setup, and building first outputs. The people who adopt fastest are those already good at writing clear briefs and organising their processes.



The Thursday AI Club

A hands-on AI club for marketers and managers. Fortnightly, 2-5pm. Led by Ann Stanley, Darren Wynn, and James Allen.

HOW EACH SESSION WORKS

Hour 1: Open Q&A - Bring any question about AI, however basic or technical. Practical answers from the team.

Hours 2-3, Track A: Workshop - Hands-on sessions where you build as you go. Prompt engineering, AI agents, Claude Code, vibe coding, and more.

Hours 2-3, Track B: Advanced - For experienced users who want individual questions answered, live software demos, and AI in action. Separate breakout for deeper work.

MEMBERS GET

Two workshops or drop-in sessions per month, access to the training portal with recordings and downloadable resources, membership of the Secret Agents community, plus a full-day hackathon every quarter.

Membership: £40 per month or £400 per year (excluding VAT)

secret-agents.ai/thursday-ai-club

7.6 Sharing What Claude Builds

Everything Claude builds lives on your laptop. That is powerful for security but creates a practical problem: how do you share your work with colleagues and clients? There are four approaches, and most teams use all of them.

GitHub as a Shared Repository

Create a shared GitHub repository that the whole team contributes to. Skills, CLAUDE.md files, templates, and project code all live here. Everyone pulls the latest version each morning. One person builds a skill, the whole team has access to it immediately. At Anicca, we built an internal skills directory with over 80 skills, all shared through GitHub so every team member can find and use them. GitHub also gives you a complete audit trail of every change.

Central Server Access (SharePoint)

If your business uses SharePoint or a shared drive, you can connect Claude to it. The safest pattern is to give Claude read access to shared folders so it can pull data and reference documents, but write access to only one designated folder. This prevents Claude from accidentally overwriting important files while still giving it access to the information it needs.

Password-Protected URLs

For sharing reports and dashboards with clients or colleagues, deploy to Netlify or Vercel. You can use Netlify Drop for a quick drag-and-drop upload, or connect through GitHub for automatic deployment. Either way, you get a shareable URL with password protection. Say "put this online with a password" and Claude does it in seconds.

Internal Skills Portal

We are building a portal that catalogues all of our skills by application and department. Staff can search for a skill, see what it does, and run it through either a simple form or a chat interface. This portal is currently under development, but it represents the direction of travel: a centralised place where every automation, report template, and workflow the team has built is discoverable and reusable.

PDFs for Email and Teams

Not everything needs to be a live URL. Claude can use Playwright to convert any HTML report into a print-ready PDF that you attach to an email or post in a Teams channel. This is often the simplest way to share a finished deliverable with someone who just needs to read it.

7.7 Governance

AI governance is not optional. Gartner predicts 40% of agentic AI projects will be cancelled due to governance failures. Here is what you need to have in place.

Your AI Policy

Every organisation using AI tools needs a written policy covering: which tools are approved, what data can and cannot be entered into AI prompts, who is responsible for reviewing AI output, and how you handle client-facing content. If you do not have one yet, the AI Policy Generator at secret-agents.ai/marketplace/vcmo-ai-policy-generator can create a first draft, though you should always have a lawyer review it before adoption.

GDPR and Data Privacy

Any prompt containing personal data sent to Claude constitutes data processing under UK GDPR. You need a lawful basis for processing and clear internal guidance on what is and is not acceptable to send. Anthropic offers a Data Processing Addendum (DPA) for API customers, where Anthropic acts as the data processor and your business is the data controller.

Key questions to address: Where is your data processed? EU data residency is available for API customers via AWS EU regions, but this is not a default setting. Check with Anthropic or your account manager. Is your data used for model training? On paid API plans and Claude Max subscriptions, Anthropic does not use your inputs or outputs to train models. If you are using Claude through the API (as Claude Code does), your data is not shared with or retained by Anthropic beyond the session. What happens with a GDPR subject access request? As the data controller, your business handles subject access requests. Anthropic, as processor, will assist under the terms of the DPA, but does not store conversation data from API usage beyond the session.

IP and Confidentiality

AI-generated content has no automatic copyright protection under current UK law – human authorship is required. Review client contracts to ensure they permit AI tool usage, and be transparent about AI-assisted outputs. If you are handling confidential client data, your AI policy must address how that data is protected during AI processing.

Human Accountability

Define who owns AI output. The AI drafts, a person signs off. If AI is 90% accurate per step, three steps without human review gives you 73% accuracy. Verification is not optional. Create override mechanisms and escalation paths for high-stakes decisions.

7.8 Running Your First AI Discovery Session

Whether you run this internally or bring in outside help, the goal is the same: identify where AI will have the most impact and build a practical roadmap to get there.

Do It Yourself

Gather your team for a half-day session. Ask: what do we spend most of our time doing? Which tasks are repetitive, data-heavy, or done the same way every time? Score each opportunity by impact and feasibility. Pick 1-2 top-scoring ideas and define a two-week pilot. Focus on the benefits quadrant: time savers and money savers first.

Bring In Experts

Anicca Digital runs structured AI Discovery Audits. This is a half-day strategic consultation where two AI strategists assess your current readiness, demo relevant AI tools using your own data, and identify the highest-impact opportunities for automation. The output is a written roadmap with prioritised recommendations, a CIO readiness score, and a compliance review. The Discovery Fee is credited against any follow-on project.

What Happens Next

SERVICE	WHAT YOU GET
AI Discovery Audit	Half-day consultation, CIO readiness assessment, written roadmap, compliance review. The fee is credited against implementation costs if you proceed.
Agent Customisation	Adapt established AI agents for your processes: content creation, social media, inbox management, diary automation. Setup from a one-off fee plus monthly maintenance per user.
Bespoke AI Build	Custom automations, voice and chat agents, product engines, lead management systems. Built on your own credentials and hosting. Development costs spread over the first year.
Training	From weekly free webinars through Thursday AI Club, to corporate AI Sprints and a full 10-day AI and Automation Course. Government-funded Skills Bootcamps also available.

"We ran this with our own leadership team and found that 60% of the time savings came from just three skills. You do not need to transform everything at once. Find the three things that hurt the most and start there."

Ann Stanley, Founder & CTO, Anicca Digital

7.9 The C10 AI Adoption Framework

Before you choose any AI tool, you need to understand where the problems actually are. The C10 framework walks through ten areas of your business – without mentioning AI at all – to identify what is broken, slow, or expensive.



The colour bands show the three phases: **Discovery** (C1-C2), **Customer & Content** (C3-C5), **Channels & Data** (C6-C8), and **Commercial** (C9-C10). Work through them in order. Only then do you choose the right tool for each problem.

START WITH PROBLEMS, NOT TOOLS

The biggest mistake businesses make is choosing an AI tool first, then looking for problems to solve with it. The C10 framework flips this – start with what is broken, slow, or expensive, then work out whether the fix is off-the-shelf software, a custom skill, or something entirely different.

CHAPTER EIGHT

08

From Personal AI to Company AI Infrastructure

Why AI departmental agents need shared skills, connected systems and change management to scale

EXECUTIVE BRIEFING

The first wave of AI adoption inside most businesses looks like a few enthusiastic individuals, or lone wolves. They install Claude Desktop or Claude Code, prove the technology works, and produce real, useful output within weeks. Other businesses arrive top-down, signing licences for Microsoft Copilot, Google Workspace AI or Salesforce Agentforce. Both routes prove value, but neither, on its own, scales.

The work changes the moment it stops being personal. A time saver for one person becomes a governance, training, knowledge management and integration challenge for the whole company.

This chapter defines AI departmental agents and AI operating systems, reviews the products already on the market, explains why neither the lone wolf route nor the bought-product route scales on its own, and lays out the four honest paths to a working AI operating system.

Claude Code is how the first wave of value is created. C⁴ is how that value becomes shared, governed and repeatable across the company.

75%

Of knowledge workers use AI at work today
Microsoft Work Trend Index, 2024

78%

Bring their own AI to work (BYOAI)
Microsoft Work Trend Index, 2024

50%+

Of UK firms now actively using AI
British Chambers of Commerce, 2026

4 to 6

Months until pioneers hit the infrastructure wall
Anicca observation, 2026

8.1 From individual tasks to whole systems

Chapters 1 to 7 of this guide showed you how to use Claude Code, Claude Desktop and the wider conversational AI coding agent ecosystem to perform individual tasks well. Each individual task tends to follow the same pattern when it is finished: defined inputs go in, an engine processes them, defined outputs come out. Anicca calls this the Christmas Cracker model.



The Christmas Cracker model captures the process for one specific task. To get the real benefits of AI we need to look at how individual tasks link together. This chapter is about how to organise individual tasks into linked workflows, so businesses can combine many crackers into a single AI departmental agent, or into a more complex set of agents that performs a whole job through many linked skills. The human interacts iteratively with a departmental manager or orchestration agent that oversees the work with the sub-agents, the individual skills or tasks underneath. The named roles look like a Recruitment Manager, an Executive Assistant, a Bookkeeper, a Customer Support Manager, a PPC Manager or a Sales Development Rep.

The term **AI employee** predates the Claude Code era. It was used by the productised platforms now mapped in section 8.4 (Sintra, Marblism, Motion) and by the orchestration-engine community around Make, n8n and Zapier, where a controlling agent at the top called a series of sub-agents underneath, each carrying out individual tasks under its management. Anicca's first AI departmental agents followed the same pattern: three n8n engines built in 2025 (Mammoth 2.1 Marketing, 2.2 Product, 2.3 AI Content Brain). The Claude Code era inherits the concept but executes it interactively.

The six working parts of any AI departmental agent

The AI itself, a model such as Claude, runs inside the skills and follows the role's instructions. The six parts below surround the AI to make it useful for a defined job.

1 Supervisor

A human owner who reviews outputs and stays accountable. In larger setups they manage an orchestration or supervisor agent that coordinates sub-agents underneath.

2 Role

The named function the agent performs, equivalent to its system prompt. The AI follows this throughout.

3 Skills / tasks

The bundle of techniques and step-by-step instructions for the job. The AI runs inside each skill, producing the output.

4 Brain, knowledge & guidelines

Templates, brand profiles, business data and rules. Tells the agent what the output should look like.

5 Tools and actions

What the agent does: writes content, sends emails, amends files, scrapes websites, reads and writes live data.

6 APIs and connections

The standard way software talks to software. Lets the agent reach your drive, inbox, CRM, ads, GA4 and other tools you already use.

At Anicca, our answer to this infrastructure gap is C⁴ (the Company Command & Control Centre). It brings together the C10 framework, shared skills, company knowledge, departmental agents and reporting dashboards in one browser-based operating system. We return to it in section 8.15.

8.2 Why personal AI does not scale on its own

Recent research from Spark Intelligence found that 83% of agency staff describe themselves as capable AI users, but only 15% are operating at the level where AI starts to change how their business competes. That gap is the gap between personal AI and company AI. The four problems below explain why personal AI cannot close it on its own.

REASON 1

Personal AI cannot be scaled to become business AI

When AI is used by early adopters alone, the value stays with the individual. The skill one person builds does not reach the team. The brand voice rules one person enforces are not the rules the others use.

The real cost is not subscriptions, which are small and visible. It is the hours every week spent reinventing work someone else has already built, while everyone who is not an early adopter waits on the sidelines.

The fix: one shared place that holds every skill, brand rule and audience profile. New staff start with a full toolkit. Improvements made by one person become available to everyone.

REASON 2

Personal AI cannot connect to the business's systems

An early adopter using Claude Desktop can draft, summarise and brainstorm. They cannot pull last month's spend from Google Ads, post a draft into WordPress, send invoices through Xero, or read this week's pipeline from HubSpot. All of that needs credentials, integrations and shared accounts the company has not given them.

So their work stays in a world of pasted screenshots and copy-and-paste. The company has someone who uses AI; the company itself is still pre-AI.

The fix: one shared place that holds API keys and access tokens centrally, with workflows that read and write from the company's actual systems on each person's behalf.

REASON 3

Personal AI improvements often stay with the individual

The early adopter improves as they learn the tool. The tool itself does not learn very much from what they have done.

When they fix a broken workflow or find a new audience profile that responds well to your campaign, the findings stay on their laptop. The next person starts from zero. After a year, the team's tools are no better than at the start. When the person leaves, their personal improvements walk out the door.

The fix: every refinement, fix and new template lives in the shared library. The tools get better as the team uses them. The company gets faster every quarter.

REASON 4

No overall view or integrated reporting across the business

A business needs more than chat. It needs dashboards that mix Google Ads, GA4, sales, finance and operations data in one place. It needs to see which AI departmental agents are running on a schedule, which tasks are stuck, what the team produced this week.

A CEO using personal AI can ask one question at a time. They cannot ask "what is the business doing across AI today?" and get an answer. Personal AI gives departments their own AI assistants; it does not give the business itself a way to see what is happening.

The fix: a single place to see the AI work of every department: dashboards, scheduled agents, today's tasks, recent outputs. One view, top to bottom.

8.3 Two ways businesses arrive at company-wide AI

Most businesses arrive at company-wide AI through one of two distinct journeys. Most businesses take one or the other. Both prove value within weeks. Neither, on its own, gets the company to real shared infrastructure.

ROUTE 1 · BOTTOM-UP

The lone-wolf route

Two or three curious people install Claude Desktop or Claude Code, work through prompts, and start producing useful output within weeks. Within two months it is on the management agenda. The pioneers prove the technology works inside the business with its own data, customers and brand voice.

The limitations: the skill lives on one laptop. Brand voice in the pioneer's head. Customer data in their email. When that person is on holiday, the skill is unusable. Within months you have four parallel versions of the same skill on four different laptops, none of them quite the same, and no clear answer to "which one is right".

ROUTE 2 · TOP-DOWN

The bought-product route

The IT director or CEO commits the company to Microsoft Copilot, Google Workspace AI, Salesforce Agentforce, ChatGPT Enterprise, Glean, or one of the productised AI employee platforms (Sintra, Marblism, Motion). Licences are issued and training is launched.

The limitations: the product was built before the conversational AI coding agent era. It was sold as "AI Copilot for your existing tools". The skills are not customised to your brand. The knowledge is not shared. The integrations stop at the vendor's product boundary. Six months later, the licence cost has not produced the change-management win that was promised at procurement.

"The question stops being 'how do we teach more people to use AI?' and becomes 'how do we give the whole company safe, shared, repeatable AI infrastructure?'"

The pivot point of every company-wide AI conversation

Both routes hit the same wall at roughly the same point. The chapter now turns to the AI departmental agents already on the market, what such an agent actually does, and the wider operating system that holds them together.

8.4 AI departmental agent products on the market

A growing market sells AI departmental agents as a productised service. The matrix below maps the six most-cited platforms against eighteen common role types, showing the named agent in each cell where it exists. Like comparing subscription tiers, but across whole AI workforces.

Sintra and Motion have the broadest role coverage; Marblism is tighter but covers the everyday-task essentials; Manus can do almost any role but in an autonomous-agent model, not a named-employee model; Genspark is multi-modal but function-based not role-based; Saner.AI is a single knowledge assistant. Based on our review at May 2026, none of the six offer brand-voice cloning at the skill level: customers of these platforms get the vendor's voice rather than their own. Vendors are moving fast, so this position may shift.

ROLE / CAPABILITY	WHAT THIS ROLE DOES	SANER.AI KNOWLEDGE WORKER	MOTION SCHEDULING	MANUS AUTONOMOUS	SINTRA BUSINESS-SCALE	MARBLISM ROLE-SPECIFIC	GENSPARK MULTI-MODAL
Executive Assistant / Admin	Calendar, email, scheduling, daily planning	Skai	Alfred	✓	Buddy / Emmie	Eva	AI Docs
Sales / SDR	Lead gen, outreach, follow-up, pipeline	—	Chip	✓	Milli	Stan	—
Marketing / Content	Campaign planning, copy, content brief drafting	—	Suki	✓	Penn	Penny	AI Slides
Customer Support	Ticket triage, response drafting, escalation	—	Clide	✓	Cassie	Cara	AI Chat
Project Management	Task tracking, deadlines, dependencies	tasks	Millie	✓	—	—	—
Recruiting / HR	Job descriptions, screening, scheduling	—	Spec	✓	Scouty	—	—
Research / Analysis	Web research, summarisation, synthesis	Smart Search	Dot	✓	Data Analyst	—	Web synth
Social Media	Content calendar, posting, engagement	—	Suki (part)	✓	Social Mgr	Sonny	AI Slides
Email Management	Inbox triage, drafting, follow-ups	Skai (part)	Alfred (part)	✓	Emmie	Eva	—
SEO / Blog Writing	Keyword research, drafts, on-page optimisation	—	—	✓	SEO Spec	Penny	AI Docs
Legal / Contracts	Contract review, clause checking, compliance	—	—	—	—	Linda	—
E-commerce	Product feeds, listings, store ops	—	—	—	E-comm Mgr	—	—
Voice / Phone Calls	Inbound calls, appointment booking, outreach	—	—	—	—	Cara	Call For Me
Image / Video Generation	Visual asset creation, ad creative, video	—	—	—	—	—	AI Image / Video
Coding / Development	Code generation, scripts, app building	—	—	✓ strong	—	—	AI Developer
Knowledge Management	Notes, search, document organisation	✓ core	AI Docs	—	Brain AI	—	—
Custom Agent Builder	Build your own role from scratch	—	✓	—	—	—	—
Brand-voice cloning	Customer-defined voice across all roles	—	—	—	—	—	—
Pricing	Entry » full plan Range across all six: Free to \$97/mo	Free \$8 / \$16	\$29 / user \$49 / user	\$20 (post-Meta)	\$39 \$97	\$22 \$39	Free \$25

Matrix reflects publicly available product information at May 2026 and is a snapshot of a fast-moving category, not a permanent market map. Its purpose is to show the shape of the category, not to rank vendors definitively.

8.5 What an AI departmental agent actually does

DEFINITION

AI departmental agents are autonomous skills that chain steps, call tools, and run on their own. They sit at the role or department level in the organisation structure: not a one-off chatbot exchange, not a tool inside someone's day, but a named role with inputs, outputs, a cadence and a human supervisor.

The clearest way to understand the difference between a single-task chatbot and an AI departmental agent is to see what the role does in a working week. Below are eight typical AI departmental agents in an SME, spanning marketing, HR, admin, customer service, sales and finance. Each shows the department it belongs to, the whole job it performs, and the bundle of skills it runs end-to-end.

ROLE	DEPARTMENT	WHOLE JOB	SKILLS BUNDLE (TECHNIQUES RUNNING TOGETHER)
PPC Manager	MARKETING	Plan and run paid search and social ads end to end	Keyword research · ad copy generation · bid management · audit reports · landing page review · competitor monitoring · budget tracking
Blog & SEO Manager	MARKETING	Produce and optimise long-form content that ranks	Keyword research · brief writing · drafting · on-page SEO · internal linking · image generation · meta and schema
Recruitment & HR Manager	HR	Source, screen, schedule and onboard candidates end to end	Job description drafting · CV screening · interview scheduling · candidate communication · reference checks · onboarding pack assembly · new-starter paperwork
Executive Assistant	ADMIN	Manage calendar, inbox, travel and meeting prep for senior leadership	Calendar management · email triage · travel booking · meeting prep notes · contact follow-ups · expense reconciliation · daily brief drafting
Customer Support Manager	CX	Triage, respond, escalate and learn from every customer interaction	Ticket categorisation · knowledge base lookup · response drafting · escalation routing · sentiment analysis · trend reporting
Sales Development Rep	SALES	Generate qualified leads, book meetings, follow up the pipeline	Lead research · prospect enrichment · outreach drafting · meeting scheduling · pipeline updates · follow-up sequences · CRM hygiene
Bookkeeper / Finance Assistant	FINANCE	Reconcile transactions, chase invoices, produce monthly close pack	Bank reconciliation · invoice matching · receipt categorisation · VAT preparation · cashflow alerts · aged debtor chasing · month-end close pack
Pricing Analyst	FINANCE	Watch competitors, analyse margins, recommend price changes	Competitor price scraping · margin analysis · stock alerts · recommendation generation · dashboard updates · daily alert email

WHY THINK IN ROLES

An AI departmental agent is mappable to a department head, a job description, a budget line. The eight roles above span six departments and could easily extend to twenty more. The role-level approach is what makes the category commercially meaningful, and what makes it possible to think about the question that opens the next section: where do these whole roles actually run?

8.6 What Anicca runs today

Each is a named role with a cadence, a brand voice, an engine, and a human supervisor. All six are operational today, not roadmap items.

Examples of Anicca's agents

AGENT	DEPARTMENT	CADENCE	GO-LIVE	OUTCOME
weekend-ai-blog	Marketing	Weekly	Apr 2026	AI in Marketing news round-up published every Sunday on anicca.co.uk and the LinkedIn newsletter
daily-pricing	Client services	Daily	May 2026 (cloud)	Charles Bentley and Books2Door pricing, stock and revenue pipeline; daily alert email
epic-monthly	Marketing	Monthly	Mar 2026	EPIC e-commerce news blog and cover image, last working day of every month
briefing-generator	Operations	Daily	Feb 2026	Daily briefing covering calendar, talks, research, todos and project status
geo-audit	New business	On demand	Mar 2026	GEO audit reports as A4 landscape decks, drawn from Waikay, GA4 and competitive data
competitor-intelligence	New business	On demand	Mar 2026	15+ section competitor intelligence reports as A4 landscape decks, used in new-business pitches and audits

Weekly Content Creator

Discovers, scrapes, drafts, illustrates and publishes a full week of marketing content for the blog, the LinkedIn carousel and the Thursday AI Club newsletter.

SKILLS BUNDLE (EIGHT SKILLS, RUN IN SEQUENCE)

news-scraper · ai-news-blog-writer · blog-image · weekly-carousel · ann-linkedin-posts · tac-newsletter-banner · wp-content-publish

- C5 Creation
- C4 Communications
- C7 Connections
- C2 Company

Cadence	Sat discovery, Sun draft, Mon publish
Outputs	~7 deliverables / week (1 blog, 1 image, 1 carousel, 1 banner, 1 newsletter, 2 LinkedIn posts)

Daily Pricing and E-commerce Checker

Pulls pricing, stock and revenue data across all e-commerce platforms, surfaces margin and stock issues, delivers a daily alert email to the pricing manager.

SKILLS BUNDLE (TEN SKILLS, RUN IN SEQUENCE)

magento-revenue · merchant-centre-scan · google-ads-pull · ga4-reconcile · apify-competitor-prices · stock-comparison · roas-by-product · daily-alert-email

- C9 Costs
- C6 Channels
- C8 Control
- C7 Connections

Cadence	Daily at 11:00 UTC (12:00 BST)
Outputs	1 daily alert email + 4 dashboard updates (pricing, e-commerce, ROAS, stock)

8.7 What is an AI operating system?

An AI operating system (AIOS) is not only a software subscription. It is the operating level that joins people, skills, data, workflows, permissions and measurement. In practical terms, an AIOS is the shared infrastructure that holds your AI departmental agents, your company knowledge and your platform connections in one place. It is not a replacement for Windows, macOS or Microsoft 365: it is the shared business level where AI skills, company knowledge, permissions, integrations, dashboards and role-level agents are managed together. AI departmental agents are the workers; the AI operating system is the workplace. One without the other either creates isolated productivity on individual laptops or unused infrastructure. The two only deliver value together.

Holds the skills

Reusable role bundles that anyone in the company can run, version-controlled and on-brand. Not stuck on individual laptops.

Holds the knowledge

Brand voice, customer data, audit findings, performance benchmarks, all in one place every workflow reads from.

Connects the platforms

One authentication system to Google Ads, Meta, HubSpot, Shopify and so on. Non-technical staff connect once and use everywhere.

The three-level hierarchy

An AIOS sits in the middle of three levels. Each level is necessary; none is sufficient on its own.

METHODOLOGY

C10

Anicca's ten-element framework for diagnosing where AI fits a business. The "what to do" level.

ARCHITECTURE

AIOS

The operating-system level that holds skills, knowledge and connections. The "how it runs" level.

ENGINE

Claude Code

Or Codex, Antigravity, Copilot, Cursor. The conversational AI coding agent that executes the work. The "what does the work" level.

This hierarchy is model-agnostic at the prompt level and engine-agnostic at the orchestration level. The operating system calls the engine: the OS is the interface a person works in, the engine is the inference and orchestration that powers the agents underneath. New engines will arrive over the next three years. The methodology and the architecture are the levels that survive an engine swap. The next page surveys the products already in market in the AIOS category, so the reader can place their own current tooling in the picture alongside the AI departmental agent products covered in section 8.4.

The AI engine that powers the system

TODAY

Claude Opus, Sonnet and Haiku via the Anthropic API. The full C⁴ portal and Anicca's six production agents all run on Anthropic models.

NEXT

Codex, Gemini and other providers can be used in place of Claude for clients whose policies or budgets need a different model. The way C⁴ co-ordinates the agents sits above the model, so changing the model only affects this one part and not the wider system.

FUTURE

Models that the company runs on its own servers, for confidential work where data cannot leave the business. Different providers chosen per agent and per task. The wider setup stays the same; only the model changes.

8.8 Examples of existing AI operating system products

A separate group of products aims at the operating system level rather than the named-role level. Most are vendor extensions of an existing enterprise tool, retrofitted for the AI era. The category is dominated by the established cloud suites; smaller specialists have grown alongside.

Microsoft Copilot + Agent 365

Microsoft 365 E7 Frontier Suite

Updated May 2026 with Agent 365 in the new Microsoft 365 E7 Frontier Suite. Copilot now manages AI agents across Microsoft 365 apps. Strong for businesses already on Microsoft; locked-in if not.

Google Workspace AI

Gemini for Workspace

Deep Gemini integration across Docs, Sheets, Gmail, Meet. Similar shape to Microsoft. Strong for businesses already standardised on Google Workspace.

Salesforce Agentforce

Salesforce platform

Agent platform for the Salesforce ecosystem. Strong if your business lives in Salesforce CRM. Locked-in if it does not.

ChatGPT Enterprise

OpenAI

Shared workspace tier with admin controls, agentic features (Operator and Agent mode) and SOC 2 compliance. Less integrated with existing business apps than the cloud-suite vendors.

Glean

Glean (independent, \$4B+)

Enterprise search plus agents. Strong for finding and summarising company knowledge across SharePoint, Drive, Slack, Notion. Limited as a general-purpose AIOS.

IBM watsonx Orchestrate

IBM

Enterprise AI agent orchestration. Strong governance and compliance posture, complex setup, expensive entry point. Aimed at large organisations.

Anthropic Claude for Enterprise

Anthropic

Anthropic's enterprise tier with Computer Use, Project Knowledge and admin controls. Best paired with Claude Code or Claude Desktop for the engine level.

Perplexity Enterprise

Perplexity

Enterprise search with sourced answers. Useful as a research component in a wider AIOS. Not a full operating system on its own.

n8n / Zapier Agents

Workflow automation

Open-source (n8n) and commercial (Zapier) workflow platforms. Powerful for connecting tools and automating flows. The original home of the "AI departmental agent" concept.

THE CATEGORY IS MOVING FAST

This list is current at May 2026 and will be out of date within months. The category lines (productivity-suite extensions, CRM-led agent platforms, enterprise search, workflow automation) are useful; the vendor list is not. Refresh quarterly.

8.9 Claude for Small Business

Anthropic launched Claude for Small Business on 13 May 2026: a packaged suite of agentic workflows and connectors, aimed at small and medium enterprises (UK term: **SME**; US term: **SME/SMB**, which gives the product its name) without an in-house developer team. It runs inside Claude Cowork on the desktop app and reaches out, through MCP connectors, into the systems the business already uses, QuickBooks for accounting, Stripe and PayPal for payments, HubSpot for CRM, Canva for design, DocuSign for contracts, Slack and Google Workspace for collaboration. The work is done *inside those systems*; Claude orchestrates and acts through them. This is a different approach from a single working canvas, where every skill, connector and dashboard sits in one place. We return to that approach in section 8.18 onwards.

WORKFLOW BUNDLES

Around fifteen ready-to-run workflows at the May 2026 release:

- **Finance:** payroll planning, month-end close, invoice chaser, margin analyser, tax-season organiser, cash-flow snapshot (30/60/90 day forecasts)
- **Sales:** lead triager, CRM cleanup, contract reviewer
- **Marketing:** campaigns and content strategist
- **Operations:** morning business brief
- **Plus:** customer service and HR helpers

Each workflow runs as a slash command, for example `/invoice-chase`, `/cash-flow-snapshot`, `/crm-cleanup`.

CONNECTOR DIRECTORY

Anthropic maintains a live connector directory at claude.com/connectors, currently listing more than one hundred integrations across finance, sales, marketing, design, productivity, data and healthcare. The directory is actively growing, and Anthropic has signalled an intention to keep broadening the partner network over time. Named partners include:

- **Finance:** PayPal, QuickBooks
- **Sales / CRM:** HubSpot
- **Design:** Canva
- **Documents:** DocuSign
- **Productivity:** Google Workspace, Microsoft 365, Slack, Google Drive

Each connector plugs Claude into the relevant system via MCP. The work then happens inside that system, with Claude acting through it rather than replicating its features inside Cowork.

GOVERNANCE & SUPPORT

- User-initiated, approval-gated actions only. No autonomous execution.
- No training on customer data on Team and Enterprise plans.
- Respects existing app permission structures.
- Free *AI Fluency for Small Business* course, co-developed with PayPal.
- Claude SMB Tour: free half-day workshops in ten US cities.
- Workday Foundation funding for fifteen solopreneur places.

Two different approaches to company AI

Claude for Small Business plugs Claude into the apps an SME already runs, with the work happening inside those apps. This is the **integration model**: keep Stripe, Canva, HubSpot and Slack, then add AI workflows that operate inside each. A **single working canvas**, by contrast, brings every skill, connector, document and dashboard together in one place, so the operator stays in one workspace and the apps become data sources rather than where the day's work happens. C⁴ takes the other approach; we describe it in section 8.18 onwards.

8.10 The technology infrastructure behind company AI

Shared AI infrastructure is a stack of nine components. Most companies have some of them already, in scattered form. Company-wide AI works when the components are connected, version-controlled and accessible to non-technical staff.

COMPONENT	WHAT IT DOES	EXAMPLE TOOLS
User interface	Where staff interact with AI	Claude Desktop, Claude Code, internal portal, forms, dashboards
Skills	Reusable instructions and workflows for each role	Skills directory, brand voice files, prompts, SOPs
Knowledge	Shared company context	SharePoint, Google Drive, internal docs, client folders
Authentication	Safe access to business tools	OAuth, Google Workspace CLI, API keys, permission controls
Repository	Stores and versions the system	GitHub, central repo, branches, pull requests
Automation	Runs repeatable actions	GitHub Actions, scheduled tasks, n8n, scripts
Database	Stores structured data and outputs	Supabase, Postgres, equivalent
Deployment	Publishes apps and dashboards	Netlify, Vercel, equivalent
Governance	Controls quality and risk	Permissions, review workflows, logs, documentation

THIS IS NOT OPTIONAL INFRASTRUCTURE

If any of these components is missing, the company has not yet moved from personal AI to company AI. It has more sophisticated personal AI. The bought-product route covers the user interface and (sometimes) the knowledge component; the lone-wolf route covers the user interface and a private skills component. Neither covers all nine.

Once the nine components are visible, the next question is no longer "what AI tool should we buy?" It becomes "who in the business owns each component, and how do we get them connected?" The next section maps the four honest paths to answering that ownership question.

WHICH OPERATING SYSTEM FOR WHICH AUDIENCE

Anicca uses Claude Desktop for knowledge work and Claude Code for development. Clients who cannot deploy either (procurement or licensing constraints) get Codex or Antigravity instead. The C4 portal sits above any of these as the operating system that holds the skills, knowledge and agents: change the engine underneath and the agents keep working.

8.11 Build it yourself, buy a system, or use a guided approach?

First useful outputs can appear in days. The full company operating system takes months.

There are four honest paths to a working AI operating system. Each suits a different business. The wrong choice wastes time and money; the right choice depends on technical confidence, headcount, and how much of the work the business wants to own.

PATH 1

DIY

One or two technical staff build the infrastructure themselves using Claude Code, GitHub, Supabase, Netlify.

Best for: technically confident teams of 5-30, with someone who genuinely wants to own the build.

PATH 2

Internal IT build

The IT department leads the project as part of normal infrastructure work, often built onto SharePoint or Google Workspace.

Best for: companies of 50+ with an existing IT function and clear governance.

PATH 3

Enterprise platform

The company commits to a vendor (Microsoft Copilot, Salesforce Agentforce, Glean) and accepts vendor-shaped workflows.

Best for: large enterprises with existing licences, strong compliance needs, low appetite for build.

PATH 4

Guided implementation

An implementation partner runs a structured discovery, builds the shared components, trains staff, and hands over. Realistic timeline is three to four months minimum for a working AI operating system; faster than the alternatives but not weeks-fast.

Best for: SMEs that want shared infrastructure but do not want to learn nine components from scratch.

DECISION FACTOR	DIY	INTERNAL IT	ENTERPRISE	GUIDED
Time to first value	2-3 months	4-6 months	3-6 months	3-4 months
Customised to brand	Yes	Yes	Limited	Yes
Internal capability built	High	High	Low	Medium-high
Vendor lock-in risk	Low	Low	High	Low
Ongoing support	Self-managed	IT team	Vendor	Partner + internal

8.12 The economics of a company AI operating system

A shared operating system, with central knowledge, shared skills and live connections to the business's systems, makes the four problems just described tractable. It is not free, and it is not effortless. Below is what it costs to do properly, and what it returns when the team gets through the transition.

WHAT IT COSTS, WHEN YOU DO IT PROPERLY

A team using AI as an operating system, building skills, running workflows and calling APIs, will need Claude Max tier accounts. A power user typically spends £80 to £160 a month per account and may max out their weekly quota. Lighter users can sit on £16 Pro or £80 Max accounts depending on their workload.

As a rule of thumb, plan for roughly £1,000 a month per 10-15 staff on Claude alone, plus the cost of the platforms the workflows connect to: Google Ads API, third-party data tools, hosting, integrations. The total is a real number, but still significantly less than employing one extra person.

WHAT IT RETURNS

Once the skills are built and your team has been trained on them, productivity per person can climb to two to five times what it was. The hard part is the transition. Old habits do not switch overnight, and there is a real dip in productivity while people learn the new way of working.

The skill that matters is fast-tracking that learning curve, getting your team through the dip in weeks rather than months. Companies that manage this see the productivity gains arrive within three to four months. Companies that do not either stall or fall back to personal AI.

The rest of this chapter sets out the roadblocks SMEs typically hit, why this becomes change management as much as AI training, how the C10 framework guides which agent to commission first, and how Anicca's own answer (C⁴) came together.

8.13 SME implementation roadblocks

Before we get to where to start, here is where this typically goes wrong. The chapter so far has been upbeat. SMEs we have spoken to have been burned, and they will recognise the patterns below. We name them so the consultant pitch on the following page lands as honest rather than evangelical.

AI operating system - SME implementation roadblocks

■ Structural roadblocks (typical change-programme failures) ■ AI-specific roadblocks (extra to a generic transformation)

ROADBLOCK	WHAT IT LOOKS LIKE IN PRACTICE	HOW TO OVERCOME IT	HOW C ⁴ SPECIFICALLY HELPS
No leadership ownership or push	AI is encouraged but not driven. No clear mandate, no accountability, no one is responsible for embedding it into workflows.	Assign explicit ownership at leadership level and tie AI adoption to operational priorities, not optional experimentation.	Leadership sees what is actually running across the business in real time through the cross-departmental dashboard. The CIO framework gives the board a structured way to diagnose where the company is strong and where it is weak. Skill-run metrics turn adoption into measurable KPIs a named owner can be held accountable for.
No protected time or investment capacity	Teams are expected to adopt AI alongside full workloads. Experimentation happens inconsistently and never stabilises into process.	Ring-fence implementation time and treat adoption as delivery work, not extra work. Accept the short-term productivity dip.	Staff start with over a hundred pre-built skills relevant to their role rather than building from scratch. Not all will apply to every business, but the ones that do produce useful output in hours rather than weeks. The Skill Creator generates a brand-specific version of a canonical role in minutes, so the time each person needs to get going is hours, not days.
Weak change management and capability	No internal champions, no shared workflows, and inconsistent usage across teams. Adoption remains fragmented and personal.	Build a lightweight change system: named champions, shared playbooks, repeatable workflows across functions.	The Skills Directory is the shared workflow manual. Named champions have a single library to point colleagues at rather than reinventing from scratch. Workflows are repeatable across functions because they live in one place, not on individual laptops.
Psychological resistance (fear and overload)	Staff fear job displacement, feel overwhelmed by constant tool and process change, or quietly revert to old methods.	Normalise AI as augmentation, not replacement. Introduce gradually, reduce cognitive load, stabilise workflows before scaling.	Brand-voice cloning means output sounds like the company, not a generic chatbot, reducing the natural "this is not how we write" pushback. The audit trail lets staff verify what AI produced and unwind it if needed. A single platform means one login and one set of habits to learn, not a dozen.
Low perceived value ("what's in it for me?")	Users do not see personal benefit, so usage stays optional. AI feels like extra effort rather than making work easier.	Anchor adoption to immediate personal wins: time saved, reduced admin, faster output. Make benefits visible at individual level first.	Pre-built skills produce useful output on the first day, not the first quarter. The dashboard shows each user how much time they have saved themselves personally, so the "what is in it for me" question has a measured answer in week one rather than month six.
Bad data foundation	Customer records split across three tools, sales pipeline in a spreadsheet, accounting categories that do not match marketing categories. AI on bad data produces output that is sort-of right but never trustworthy.	Audit and consolidate the data foundation before scaling AI on top. Designate a single source of truth per data domain.	Shared authentication and a shared company knowledge base mean data converges in one place every workflow reads from. C ⁴ does not fix the legacy mess on its own, but it stops adding new fragmentation and highlights the gaps that need closing first.
No measurement, no ROI proof	Board approves AI budget; six months later there are anecdotes but no numbers. Initiative dies of accounting reasons rather than technical ones.	Set up a measurement framework on day one. Track time saved per agent, runs per department, output quality. Report monthly.	Run history and skill-level metrics are built in from the start. Time saved per agent, runs per department, success rate per workflow are all recorded automatically. Measurement is part of the platform, not something to be added on afterwards.
Vendor or model lock-in fear	Boards worry about committing to one platform (Claude, OpenAI) only to have prices spike or features removed. Hesitation around long-term commitment.	Choose a platform that abstracts the underlying model so the engine can be swapped without rewiring workflows.	C ⁴ is model-agnostic and engine-agnostic. Switch the engine underneath, whether Claude, OpenAI, Gemini or a self-hosted model, and the agents above keep working. Removes the vendor lock-in question that worries boards.

For a fuller view of what individual skills look like and how a team builds its first ones, see Chapter 5 (Skills), particularly section 5.8 Your First Week Plan and section 5.9 What to Build Next.

8.14 Why this becomes change management, not just AI training

The shift from personal AI to company AI is not really a technology project. It is a change-management project with technology inside it. Most SMEs underestimate this, and that is where bought-product roll-outs stall.

What AI training covers

- How to write a useful prompt
- How to structure a Claude Code session
- How to use Claude Desktop, Cursor, Copilot
- How to think about hallucinations and accuracy
- How to install and run a single skill

What change management covers

- Which roles should be shared, and which should stay personal
- Who owns the skills directory, the knowledge base, the authentication system
- How outputs get reviewed and signed off
- How non-technical staff are trained, supported and trusted
- How the system is maintained when the original pioneer is not in the room

Both are needed. Training without change management leaves you with a smarter version of the lone-wolf problem: more pioneers, same fragmentation. Change management without training leaves you with a beautifully governed system that nobody knows how to use.

"Many organisations can start this themselves. In fact, they should. But once AI moves from personal productivity into shared infrastructure, most SMEs need a structured implementation process."

The honest consultant pitch

A good implementation partner helps map the use cases, design the repository structure, set up safe authentication, create shared skills, connect existing knowledge sources, define governance, train staff, and leave the business with a system it can continue using without relying on one or two internal champions. That is the work. It is not glamorous, and it is not the same as installing an AI tool.

What this looks like when it works

"Working with Anicca gave us a bespoke CustomGPT for creating marketing calendars, writing blogs and other content trained on our brand guidelines. We also used their Mammoth Content Creator for generating product descriptions across each category of our website. We are now creating great content in a fraction of the time it used to take."

Charlie Hubbert · Ecommerce Executive, Charles Bentley

"I joined the 2-day AI Sprint and it was a great choice. I am looking to adopt AI Agents and automation in my every day work and the Anicca workshop has provided me with the best tool insights and easy to follow steps to implementation. I now have a lot of planning to go and do but this course has been great, thanks guys!"

Jo-Anne Cluer · Clarendon Speciality Fasteners

8.15 The C10 framework as the diagnostic

Once you accept that company-wide AI is an infrastructure project, the next question is where to start. The C10 framework walks ten dimensions of the business and surfaces, for each one, the gap that an AI departmental agent could close. It is the diagnostic that turns "we should do something with AI" into a prioritised list of named roles to commission first.

C1 Challenges Discovery agent	C2 Company Daily Brief Writer	C3 Customers Customer Health Reporter	C4 Communications Email Triager	C5 Creation Content Engine
C6 Channels PPC Manager	C7 Connections Integration Manager	C8 Control KPI Summariser	C9 Costs Pricing Analyst	C10 Compliance Policy Reviewer

How C10 maps to the C-suite

ROLE	C-ELEMENTS THEY TYPICALLY OWN	FIRST AI DEPARTMENTAL AGENTS THAT BENEFIT THEM
CEO	C1 Challenges, C10 Compliance	Discovery Agent, Daily Briefing, AI Policy Manager
CMO	C3 Customers, C4 Communications, C5 Creation, C6 Channels	Customer Avatar Council, Content Engine, PPC Manager, Blog & SEO Manager
COO	C2 Company, C7 Connections	Daily Brief Writer, Integration Manager, Executive Assistant
CTO	C7 Connections, C8 Control, C10 Compliance	AI Policy Manager, Integration Manager, KPI Summariser
CFO	C9 Costs, C10 Compliance	Bookkeeper / Finance Assistant, Pricing Analyst, AI Policy Manager

The diagnostic is the same whichever path the company chooses, DIY, internal IT, enterprise platform, or guided implementation. C10 tells you **what** to build first; the path you choose in section 8.10 (Build / Buy / Guided) tells you **how**.



8.16 Introducing C⁴: Company Command & Control Centre

C⁴ is the AI operating system Anicca built and runs day to day. A single working canvas where AI skills, dashboards, shared knowledge and connections to existing systems all come together for the whole company.

It sits in the AIOS category described in section 8.7, alongside the enterprise platforms in section 8.8 (Microsoft Copilot, Google Workspace AI, Salesforce Agentforce, IBM watsonx) and the SMB integration suite in section 8.9 (Claude for Small Business). C⁴ takes the third approach: a single working canvas where the company's AI work is done, rather than the add-on or integration approaches taken by the others.

For a business moving from one or two AI users to AI used across the whole company, C⁴ provides a single place to bring that work together. Skills, dashboards, shared knowledge base and connections to the systems the business already runs are all configured around the business's own structure, not generic templates. The internal version runs today across the Anicca group; the customer-facing SaaS portal is in late-stage prototype and being built through 2026. The architecture described in the three pieces below has been proven inside the Anicca group on real workflows.

The three architectural pieces that make C⁴ work

Three pieces turn the AIOS principles into a working product. Each one addresses a specific implementation roadblock from section 8.13, shown in the “Solves” line under each card.

Anicca Skill Creator Brand-voice cloning at the role level. Take one canonical role (a Blog and SEO Manager, an Audit Report Builder) and apply a brand profile to get back a brand-specific version in minutes. Staff start with over a hundred pre-built skills relevant to their role rather than building from scratch. SOLVES: no protected time, weak change management, low perceived value.	Your company knowledge in one place Brand voice, ICPs, product catalogue, customer data, audit findings and operating procedures all live in one place every workflow reads from. Output stays on-brand and on-fact, because the knowledge is shared, not re-described in every prompt. SOLVES: bad data foundation, psychological resistance to off-brand output, weak change management.	Integration Manager Connects C⁴ to your company's existing software (Google Ads, HubSpot, Xero, Magento, Shopify and more) so agents can read and write directly to those systems. A single place to set up, monitor and renew the connections, with a record of every action taken. SOLVES: no technical setup, no measurement and no proof of value, vendor lock-in fear.
--	---	---

A note on API access. Company-wide platforms used by the whole team, where everyone connects through one shared business account, are relatively simple to link, and we set them up during onboarding. Individual accounts where every member of staff has their own login, such as Gmail and Microsoft Teams, need a different approach. This is currently under development at Anicca.

8.17 How C⁴ came together

The following section is not a generic market overview. It explains how Anicca has applied the AIOS principles in its own working environment, and how that work has developed into C⁴.

C⁴ did not start as a single product idea. It is the convergence of five parallel threads that Anicca built or validated independently between 2023 and 2026. Each thread addresses one piece of the company-level AI gap described earlier in this chapter (section 8.2), and together they form one product. C⁴ gives leaders control without taking autonomy away from teams: shared skills, shared knowledge, shared connections and shared visibility.

THREAD 1 · FOUNDATIONAL The C10 framework Ten organisational elements covering every knowledge-work function, from C1 Challenges through C10 Compliance. Maps a company's work onto a shared vocabulary every department can diagnose against. NAVIGATION, DIAGNOSIS AND DISCOVERY	THREAD 2 · LAUNCHED MAY 2025 secret-agents.ai Anicca's first prototype agent storefront, built on n8n and released externally in May 2025. Twenty-five-plus ready-made agents that any business could pick up. Many of those agents are now candidates for conversion into skills in the Anicca Skills Directory. ANICCA'S PROTOTYPE AGENT MARKETPLACE	THREAD 3 · SHOPIFY APP JAN 2026 Armadello Analytics App Started as a Shopify analytics app on Zoho Reports. The latest version uses the dashboard infrastructure behind the Anicca client portal and C⁴, packaged as a standalone reporting portal for e-commerce clients and a Shopify app embedded in the merchant's admin. E-COMMERCE ANALYTICS PRODUCT	THREAD 4 · BUILT SINCE FEB 2026 The Anicca Skills Directory Anicca's first attempt at a common portal for housing every skill, prompt, role and template, so any team member's work can be reused by everyone. Over a hundred skills built so far, mapped against the C10 framework. Designed to merge with secret-agents.ai over time. ANICCA'S FIRST SKILLS HOUSING PORTAL	THREAD 5 · BUILT SINCE FEB 2026 The Anicca client portal Audit-insights portal built first for live clients, multi-tenant on Supabase, eight production dashboards across Google Ads, GA4, Magento and HubSpot. The dashboard infrastructure now powers Armadello. DASHBOARDS AND CLIENT DATA
---	---	--	---	--

C⁴: Company Command & Control Centre

The five threads converged in April 2026 as the **A-AIOS** (Anicca AI Operating System), the custom-assembled system Anicca runs day to day. In May 2026 the productised version was named **C⁴**: Company Command & Control Centre; one browser-based product, owned by the business, combining control of the framework, the agents, the skills and the dashboards. Cross-departmental, model-agnostic, engine-agnostic, configurable per role. The internal version runs today across the Anicca group; the customer-facing SaaS portal is in late-stage prototype and being built through 2026.

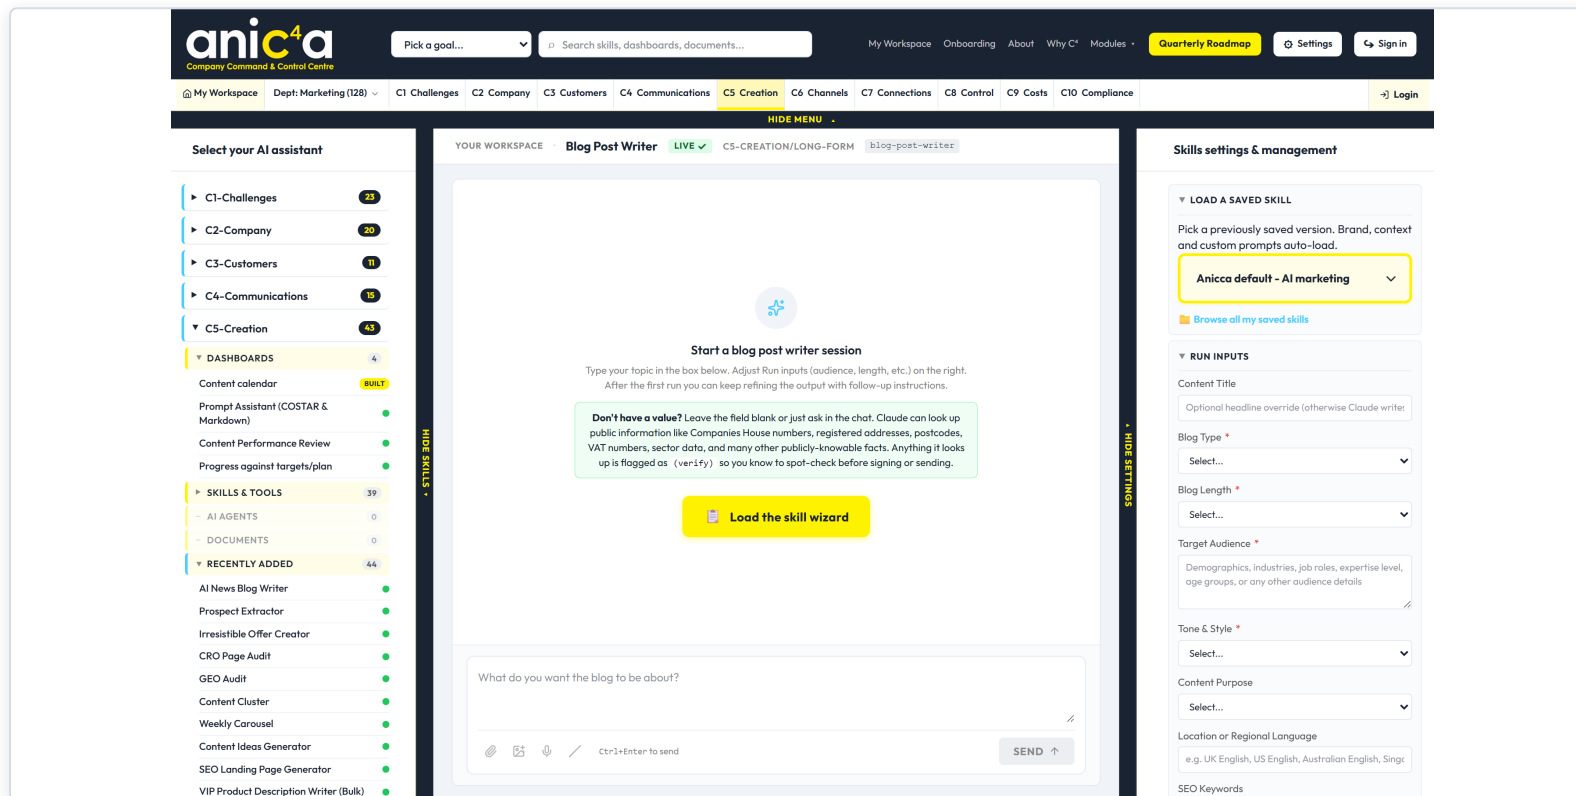
C⁴ runs today across Anicca Digital, Anicca Insights and Armadello, built on Claude Code, GitHub, Supabase, Netlify and SharePoint. The productised SaaS version is in development; the internal version is what we use day to day.

Why it works

Closes the pioneer asymmetry Skills live in the shared Skills Directory, not on individual laptops, so the value rolls out to the whole team.	Unblocks system access The single authentication in A-AIOS lets non-technical staff connect to Google Ads, HubSpot, Xero and other platforms without IT setup.	Compounds over time Every refinement one person makes becomes available to everyone via the Skills Directory and the Skill Creator.	Gives the business one view The client portal mixes data across departments and clients so the CEO can see what is happening in real time.	Solves the UK paralysis The C10 framework lets buyers diagnose their AI readiness, then onboard in weeks rather than years.
---	--	---	--	---

8.18 What C⁴ looks like

The earlier sections of this chapter described two existing routes for getting AI into a company. The enterprise platforms in section 8.8 (Microsoft Copilot, Google Workspace AI, Salesforce Agentforce and the like) are aimed at large organisations with in-house IT teams. The integration suite in section 8.9 (Claude for Small Business) is aimed at small businesses, with AI reaching out into the apps they already use. C⁴ is built differently from both. The image below shows the live version of C⁴ that the Anicca team uses day to day. The next section, 8.19, explains how it is built, and why those choices matter for a business that is moving from one or two AI users to using AI across the whole company.



The C⁴ working canvas. Left: the department picker and per-department skill tree. Centre: the chat-style work area, here showing a blog post writer in C5 Creation. Right: the saved settings and run inputs for that skill. Each user opens C⁴ to their own workspace with their own skills. The shared brand voice, customer data and operating procedures are configured separately and read into every workflow.

8.19 What is different about C⁴

There are six practical differences between C⁴ and the products described in sections 8.8 and 8.9. Each one is designed to address a problem that comes up repeatedly when a business tries to take AI from one or two enthusiastic users to the whole company.

1. Set up around the business before anyone starts using it.

Most AI platforms start every user with a blank prompt. C⁴ starts at the company level. A structured discovery (the C10 framework described in section 8.15) captures the company's departments, workflows, roles, brand voice and operating procedures first. The workspace is then configured around what was captured, so when staff log in, the AI already understands their business.

2. Each user has their own workspace with their own skills.

Every staff member opens C⁴ to a workspace built for their role: the skills they use, the dashboards they need, the permissions they hold. The marketing team and the operations team see different things on screen, but they all draw on the same shared brand voice, customer information, audit findings and operating procedures, so AI output stays consistent across the business.

3. Forms for the repeatable, chat for the rest.

Most company tasks are not free-form conversations. They are 'produce this report' or 'write this email'. C⁴ uses structured forms for the repeatable workflows, with the user filling in the inputs (audience, length, tone, examples) once rather than re-explaining each time. The chat interface stays there for refinement, follow-up and one-offs, so the same workspace handles both styles of work.

4. Connected to the company's software and the user's own accounts.

Most businesses already pay for Google Ads, HubSpot, Xero, Shopify, Microsoft 365 and similar systems. C⁴ connects at the company level through APIs, and where appropriate can connect to user-level accounts such as email, calendar, chat and drive. Skills read this data automatically through one shared authentication, so users do not have to log in to each system separately or copy and paste between five tabs to do their work.

5. The AI that powers C⁴ can be swapped without breaking the workspace.

C⁴ is built using Claude today, but the workspace, the skills, the forms and the shared knowledge are not tied to Claude. If a business prefers OpenAI, Google's Gemini or any future model that becomes available, the AI underneath can be swapped while everything above stays in place. The business is never locked in to one vendor.

6. A guided onboarding pathway, not just software access.

Setting up a company-wide AI workspace involves more than logging in. There is infrastructure to configure: code repository, shared knowledge storage, hosting, shared authentication, brand voice and operating procedures. C⁴ includes a structured onboarding process that handles all of this, delivered today by Anicca or a certified partner, and in time as a guided self-serve route through the portal.

The AI models powering C⁴ today, Claude, GPT and Gemini, will keep improving and getting cheaper over the next few years. The model on its own is not what makes C⁴ different. What matters is everything around it: the discovery, the per-role workspaces, the form and chat interface, the connections to every system the business already uses, and the way AI is adopted across the whole company. That is the part that takes time and effort to build, and it is what C⁴ provides as a working product.

8.20 Where to start

Where to start: three free assessments on secret-agents.ai

ROI calculator

Estimate what an AI operating system could deliver for a business your size, in time saved, costs reduced, and revenue uplift.

secret-agents.ai/calculators

Readiness quiz

A 60-second self-assessment that tells you whether your business is ready to adopt company-wide AI, or whether foundations need fixing first.

secret-agents.ai/quizzes

Free C10 audit

A 10-minute self-audit that scores your business across all ten C-elements and produces a personalised priority report.

secret-agents.ai/c10-audit

If you want a conversation

Free 30-minute discovery call

Walk through your audit results with us. No pricing, no slides, no pressure. Just your situation and an honest read on the best next step.

Free webinar: AI operating systems for SMEs

A live walk-through of what implementation actually looks like, covering the nine-component infrastructure and the four routes to get there. Run regularly; check the link below for the next date.

The next three years of agentic AI will be defined by which companies built shared infrastructure and which kept patching personal AI. Because the C10 framework is engine-agnostic, it continues to apply no matter which model or engine arrives next. Build it now, while the cost of starting is still low.

Glossary (A–G)

AEO (Answer Engine Optimisation) – See GEO. An alternative name for the same practice of optimising content to appear in AI-generated answers

Agentic AI – AI that accomplishes goals with limited supervision, planning steps and adapting to obstacles

AI Departmental Agent – An autonomous skill, or chain of skills, that performs a whole role-level job. Sits at the role or department level with inputs, outputs, a cadence and a human supervisor

AIO (AI Overviews) – Google's AI-generated answers that appear at the top of search results, summarising information from multiple sources

AIOS (AI Operating System) – Shared infrastructure holding an organisation's AI departmental agents, company knowledge and platform connections in one place. Without it, agents are stranded on individual laptops

Anthropic – The AI company that makes Claude, founded by former OpenAI researchers

API (Application Programming Interface) – A way for software to talk to other software. Claude uses APIs to pull data from Google Ads, Analytics, and other services

Artifact – Content Claude creates in the web interface (document, code, table) that can be published as a web page

Authentication – Verifying a user's identity before granting access. OAuth tokens, API keys, and passwords are all forms of authentication

Auto Mode – Permission mode where an AI classifier approves routine actions automatically while blocking risky ones

Base64 – Encoding images as text so they embed directly inside documents, ensuring colours and proportions stay accurate every time

Brand Bible – Rules defining your visual identity (fonts, colours, logos, tone) that Claude follows across all output

Brand Voice – A skill capturing your personal or company writing style with examples and rules, so Claude sounds like you in every output

Browser Automation – Claude controlling a real web browser (via Playwright or Chrome MCP) to navigate sites, capture screenshots, fill forms, and extract data where no API exists

C⁴ (c4) Portal – Anicca's AI operating system: a web app that holds the skills library, the company knowledge base and the department-level agents in one place, with a multi-departmental view mapped to the C10 framework

Central Knowledge Base – The company library inside C⁴ that holds tone, brand profiles, ICPs, product catalogues and operating procedures, so every departmental agent reads from one source instead of starting from scratch each time

ChatGPT – OpenAI's AI chatbot, the most widely known AI assistant. Claude is made by a different company (Anthropic)

Chrome MCP – Connects Claude to your real Chrome browser so it can interact with sites using your logged-in session. Used for dashboards and tools that require authentication

Claude Code – Anthropic's AI-powered operating environment that reads your files, remembers your preferences, runs skills, and produces finished work on your machine

Claude Cowork – The mode inside Claude Desktop where Claude can read and write files in a folder you select. Sits alongside Chat and Code as one of three Claude Desktop modes

Claude Projects – Workspaces on claude.ai where you upload reference files and set custom instructions. The web-based predecessor to Claude Code

CLAUDE.md – A text file of instructions placed in your project folder. Claude reads it automatically at the start of every conversation

CLI (Command Line Interface) – A text-based way to interact with software by typing commands instead of clicking buttons. Claude Code started as a CLI tool

Codex – OpenAI's cloud-based coding agent. Runs tasks in a sandboxed environment, similar to Claude Code but without local file access

Computer Use – AI that can see, click, and navigate your screen like a human user. Claude, Manus, and Google Mariner all offer versions of this

Context Window – The amount of information Claude can hold during a session. Like a desk, only so much fits before you need to tidy up

Copilot Cowork – Anicca shorthand for the Microsoft 365 Copilot mode where Copilot operates on files inside your M365 cloud tenant. Microsoft's equivalent to Claude Cowork, built primarily on OpenAI GPT models

COSTAR – A prompting framework: Context (background), Objective (what you want), Style (tone), Tone (emotion), Audience (who reads it), Response (format)

Credential – A secret value (API key, password, token) that proves your identity to a service. Store in .env files, never in code

CRO (Conversion Rate Optimisation) – Improving a website so more visitors take the desired action, such as buying, signing up, or enquiring

CSS (Cascading Style Sheets) – The code that controls how web pages look, including colours, fonts, spacing, and layout

CSV (Comma-Separated Values) – A simple spreadsheet format that most tools can import and export. Claude analyses CSVs for reporting

Cursor – A code editor built on VS Code with deep AI integration. Supports Claude Code as an extension

Custom Instructions – Rules and preferences you set for a Claude Project on the web, equivalent of CLAUDE.md for web users

Deploy – Publishing a file or application to the web so others can access it via a URL

Department Matrix – The two-axis map of departments (X-axis) against C-elements (Y-axis) that makes the C⁴ portal navigable: lets a marketing or finance user find the skills relevant to their role without trawling the whole library

.env (Environment File) – A hidden file storing credentials and configuration values. Claude reads these to connect to services like Google Ads or Supabase

Engine (AI sense) – The inference and orchestration that runs underneath an AI operating system: the Anthropic API, Claude Agent SDK, GitHub Actions and similar offerings from other providers. The OS calls the engine

EPUB – An ebook format readable on Kindle, tablets, and phones. Claude can convert documents to EPUB

Extension – A plugin installed inside Cursor or VS Code to add features. Claude Code is installed as an extension

GEO (Generative Engine Optimisation) – Optimising content so it appears in AI-generated search results from ChatGPT, Perplexity, Google AI Overviews, and similar tools

Git – Version control software that tracks every change to your files. GitHub is the cloud service built on top of Git

GitHub – Cloud service for file backup with version history, like OneDrive but keeps every version so you can always go back

Google Antigravity – Google's browser-based AI coding environment. An IDE that runs entirely in your browser with Gemini built in

GPTs – OpenAI's shareable prompts with uploaded files, available in ChatGPT

Guardrails – Rules and limits that prevent AI from doing things it should not, such as deny lists, content filters, and permission modes

Glossary (H–Z)

Hallucination – When AI generates confident but incorrect information. Always verify facts, figures, and claims in AI output before using them

Handover Document – A file recording project progress, decisions, and next steps. Maintains context between sessions

Hook – A command that runs automatically before or after Claude performs an action, such as formatting code before every save

HTML (HyperText Markup Language) – The language used to create web pages. Claude produces pixel-perfect HTML that converts to PDF

Human-in-the-Loop – A workflow where a person reviews and approves AI output before it is used. The default approach for anything client-facing

ICP (Ideal Customer Profile) – A description of the type of company or person most likely to buy from you, used to focus marketing effort

IDE (Integrated Development Environment) – The tool developers use to write, test and run software, like a word processor for code. Examples: Cursor, VS Code, Google Antigravity, JetBrains. Claude Desktop is not an IDE – it is a desktop app with a built-in Code tab

Integration Manager – A departmental agent that sets up and maintains the API connections between C⁴ (or any AI operating system) and external tools such as Google Ads, GA4, the CRM and the e-commerce platform

JSON (JavaScript Object Notation) – A structured text format for storing data and settings. Configuration files like settings.json use this format

Knowledge File – A document uploaded to a Claude Project for reference, such as brand guides or data files

LLM (Large Language Model) – The type of AI that powers Claude, ChatGPT, and Gemini. Trained on vast amounts of text to understand and generate language

Manus – A cloud-based AI agent that controls a virtual computer to complete tasks. Competes with Claude's Computer Use feature

Markdown (.md) – Simple text format using symbols for formatting (# headings, ** bold, – bullets)

Marketplace – In Claude Code: a searchable directory of pre-built skills and integrations you can install with one command. Separately, the term also refers to the VS Code / Cursor extension marketplace where you install the Claude Code extension itself

MCP (Model Context Protocol) – Open standard connecting AI tools to external services like Jira, Slack, Google Sheets, and browsers

MEMORY.md – Persistent file storing project status across conversations, your project dashboard

Metadata – Name and description at the top of a SKILL.md file. Tells Claude when to use the skill

n8n – Self-hosted workflow automation tool. Good for high-volume, rule-based integrations

Netlify – A web hosting platform that turns HTML files into live, password-protected websites with a shareable URL

OAuth – A secure way to grant apps limited access to your accounts (such as Google or Microsoft) without sharing your password

Operating System (AI sense) – The place a person works with AI: the C⁴ portal, Claude Desktop, Claude Code, Cursor, Codex or Antigravity. The interface sits above the engine; the engine does the inference, the operating system holds the workflows, skills and knowledge

Opus – Anthropic's most capable model line for complex reasoning, coding, agentic workflows and long-context work. Check the model selector in your Claude app for the current default

Orchestration – Coordinating multiple AI agents to work together on a task, with one agent managing the others

PDF (Portable Document Format) – A fixed-layout document format. Claude generates PDFs from HTML using Playwright for client deliverables

Playwright – An open-source browser automation tool. Claude uses it to take screenshots, generate PDFs, and interact with websites programmatically

PowerShell – Windows command-line tool, similar to Terminal on Mac. Claude Code can run commands through it

PPC (Pay-Per-Click) – Paid advertising where you pay each time someone clicks your ad, typically on Google or social media

Prompt – The text you type to tell AI what you want, whether a question, instruction, or description

RAG (Retrieval-Augmented Generation) – Giving AI access to your own documents so it can answer questions using your data, not just its training

Reference Files – Supporting documents inside a skill folder (examples, style guides, seed images)

Repository (Repo) – A project folder stored on GitHub containing all your files, their history, and every change ever made

ROAS (Return on Ad Spend) – Revenue generated for every pound spent on advertising. A ROAS of 5 means five pounds back for every one spent

Seed Image – A visual reference in a skill folder so Claude can see what output should look like

SEO (Search Engine Optimisation) – Improving a website so it ranks higher in Google and other search engines for relevant searches

settings.json – Configuration file for Claude Code permissions and allowed commands

Skill – Reusable instructions stored as files, triggered by slash commands. The building block of reproducible AI

Skill Creator – A skill inside C⁴ that lets a non-developer turn a repeatable task into a named departmental agent without writing code

Skills Directory – The browsable catalogue of skills inside C⁴, organised by C-element and by department, so a user can find the right skill for the job in one click

SKILL.md – The main instruction file inside a skill folder, plain Markdown with a name and description

Slash Command – A shortcut (like /carousel or /proposal) that triggers a skill in Claude Code

SQL (Structured Query Language) – The language used to query and update databases. Claude writes SQL to pull data from Supabase and other database systems

Supabase – A cloud database platform. Claude uses it to build applications with user logins, data storage, and real-time updates

Terminal – A plain text window that runs commands. One of the three Setups (or Wrappers) for Claude Code, alongside Claude Desktop and IDEs like Cursor / VS Code. Primarily used by developers – less suitable for marketers and business owners who need visual tools to create, preview and review documents

Token – The unit AI uses to measure text, roughly one per word. Your context window has a token limit

Vercel – A web hosting platform similar to Netlify, used to deploy applications and reports to live URLs

VS Code – Free code editor by Microsoft. Claude Code installs as an extension

Wispr Flow – A paid voice-to-text tool (\$8-13/month) that gives you one hotkey across every app. Useful if you work in Terminal, or if you want consistent voice input everywhere. Claude Desktop Cowork, Cursor and the Desktop Code tab all have voice built in already

Appendix A: Setting Up Your Permissions

Permissions live in a file called `settings.json` inside your `~/ .claude/` folder. Three rule types control what Claude can do:

RULE	WHAT IT DOES
Allow	Claude does this automatically, no prompt needed
Ask	Claude always asks you first before proceeding
Deny	Always blocked. Claude cannot do it even if you approve

Deny always wins. If something is denied at any level, nothing can override it.

Recommended Starter Settings

Create this file at `~/ .claude/settings.json`. This is a safe starting point - you can add more commands to the allow list as you get comfortable.

Option 1: Auto Mode (Recommended for Claude Desktop)

Auto mode uses an AI classifier to approve routine actions automatically while blocking anything risky. This is the recommended way to use Claude Code if you are using the Desktop app.

In Claude Desktop, toggle **Auto accept edits** at the bottom left of the Code tab. That is all you need to do.

The classifier decides what is safe and what needs your approval. You can still add deny rules in `settings.json` to block actions that should never happen (like force-push or recursive delete), because deny always wins regardless of mode.

Setting Up Bespoke Permissions

Option 2: Bespoke Permissions (Cursor / VS Code)

In Cursor or VS Code, you can enter auto mode for the current session by pressing **Shift+Tab** to cycle through permission modes. For persistent fine-grained control across sessions, write explicit allow, ask, and deny rules in a settings file. This is effectively what auto mode does behind the scenes – you are setting the rules yourself.

SETTINGS.JSON - BESPOKE CONFIG

```
{
  "permissions": {
    "allow": [
      "Read", "Glob", "Grep",
      "Bash(git status)", "Bash(git add *)",
      "Bash(git commit *)", "Bash(git pull)",
      "Bash(git log *)", "Bash(git diff *)",
      "Bash(python *)", "Bash(node *)",
      "Bash(ls *)", "Bash(pwd)", "Bash(date *)"
    ],
    "ask": [
      "Edit", "Write", "Bash",
      "WebFetch", "WebSearch"
    ],
    "deny": [
      "Bash(git push --force*)",
      "Bash(rm -rf *)"
    ]
  }
}
```

This auto-approves reading and safe commands, asks before editing or running unknown commands, and permanently blocks dangerous operations. Type `/permissions` to see your active rules at any time.

Appendix B: What to Say to Claude

The most common question from new users: "What do I actually type?" Here are the exact phrases that work.

Starting a Session

SITUATION	WHAT TO SAY
Picking up a project	<i>"I am working on the Cloudfy audit. Read the handover doc."</i>
Starting something new	<i>"Build a competitor report for [company]. Read the competitor-intel skill."</i>
Quick one-off task	<i>"Summarise this PDF and pull out key action items."</i>

During a Session

SITUATION	WHAT TO SAY
Save progress	<i>"Save progress to handover - finished sections 1-3."</i>
QC the output	<i>"Compare this to the reference and list every difference."</i>
Fix something	<i>"The logo is too small on page 3. Make it 20% larger."</i>

Ending a Session

SITUATION	WHAT TO SAY
Summarise and save	<i>"Write a handover for this session."</i>
Save and back up	<i>"Commit and push to GitHub."</i>

Quick Setup Tips

For Cursor / VS Code users. Claude Desktop is the recommended starting point for most readers; the shortcuts and theme tips below apply if you are using the IDE route.

Switch to Light Mode

Press **Ctrl+Shift+P**, type "Color Theme", choose a light theme. Claude Code inherits your VS Code theme automatically.

Change Permission Mode

Press **Shift+Tab** to cycle through modes: Default, Accept Edits, Plan, Auto. Or click the mode label at the bottom of the Claude panel.

Check Your Permissions

Type **/permissions** in any Claude Code session to see all your active allow, ask, and deny rules.

Keyboard Shortcuts

SHORTCUT	WHAT IT DOES
Shift+Tab	Cycle through permission modes (Default, Accept Edits, Plan, Auto)
Ctrl+Shift+P	Open command palette (search for Claude commands)
Ctrl+Shift+X	Open extensions panel
Escape	Cancel Claude's current task

TEAM SETUP GUIDE

For detailed onboarding including GitHub setup, branch management, and daily automation scripts, see the **Training Portal** (Module 3: Getting Set Up) or contact your team lead.

Your Next Step: The Thursday AI Club

This guide has given you the strategic case and the core concepts. The **Thursday AI Club** is where you put it into practice alongside other business leaders. Membership includes fortnightly live workshops, a quarterly hackathon, and full access to our community portal with Claude Code training, Q&A recordings, downloadable skills, and ongoing support as the tools evolve. The **Claude Code training section** of the portal is also open to all guide readers, with step-by-step setup guides and video walkthroughs.

OPEN ACCESS

For Your Team

- Installing Claude Code (step by step)
- Permission configuration and starter settings
- Session phrases - what to type and when

OPEN ACCESS

Video Walkthroughs

- Building your first CLAUDE.md
- Creating and testing a skill
- Memory, handovers, and session management

MEMBERS ONLY

Thursday AI Club Members

- Full portal access included with membership
- Fortnightly live workshops
- Q&A recordings and community discussion



Join the Thursday AI Club

Fortnightly workshops, a community portal with recordings and resources, the Claude Code training library, and a quarterly hackathon. Everything you need to learn with people who are doing it alongside you, included with membership.

ann@anicca.co.uk • secret-agents.ai/thursday-ai-club

Appendix C: References & Further Reading

Sources for the May 2026 update

- **Claude Design** – Anthropic Labs announcement, May 2026. See anthropic.com/news
- **Claude for Small Business** – Anthropic product launch, May 2026. Connectors documentation at docs.claude.com
- **Claude Opus 4.7** – Anthropic model release notes, May 2026. See docs.claude.com/en/docs/about-claude/models
- **/effort xhigh** – Claude Code reasoning levels documentation, docs.claude.com/en/docs/claude-code
- **Claude Code routes (web, JetBrains, Slack, CI/CD)** – Anthropic platform documentation, May 2026
- **Claude Code memory behaviour** – CLAUDE.md and remembered preferences documented at docs.claude.com/en/docs/claude-code
- **Claude Code permissions** – settings.json reference at docs.claude.com/en/docs/claude-code
- **AI operating systems category** – industry coverage from a16z, Foundation Capital and Bessemer Venture Partners (2025-2026)
- **McKinsey, “The state of AI in early 2025”** – mckinsey.com
- **Microsoft Work Trend Index 2024** – microsoft.com/worklab/work-trend-index
- **Adobe State of Creativity 2026** – adobe.com
- **DSIT UK Business AI Survey 2025** – Department for Science, Innovation and Technology, gov.uk
- **British Chambers of Commerce Quarterly Economic Survey 2026** – britishchambers.org.uk

The sources that informed Chapter 8 and the wider C⁴ thinking behind this guide.

From Anicca

C⁴ Strategy and Market Research, v2.4 and v2.5 (May 2026). The detailed strategy and market research that sit behind this chapter.

External sources

- **Stu Long**, how-to-ai.guide – voice-cloning workflow.
- **Nate Herk**, github.com/nateherkai/AIS-OS and the AI Automation Society community on Skool – personal-AIOS starter kit.
- **Liam Ottley**, morningside.ai – productised AIOS sales training.
- **Dave Ebbelaar**, datalumina.com – engineering perspective on AIOS architecture.
- **Saner.AI**, "AI Employees: We Tested The 6 Best Platforms" – blog review covering Sintra, Marblism, Motion, Manus, Genspark and Saner.AI itself.

Ready to Start?

This guide has given you the strategic overview and the practical steps. Here are several ways to move from understanding to implementation.

Join Thursday AI Club

Fortnightly workshops where agency owners and business leaders explore AI tools hands-on. No theory - just practical implementation with peers. Bring your laptop, bring your questions.

Book a Strategy Call

A focused conversation about where agentic AI fits in your organisation. We will identify your highest-value use cases and give you a practical starting plan.

"The future is already here - it is just unevenly distributed."

William Gibson

anicca.co.uk • ann@anicca.co.uk

READ THE BLOG

Further Reading on anicca.co.uk/blog

- [From Claude Projects to Claude Code](#) - A beginner's guide to making the switch
- [5 Claude Code Skills Every Business Should Build First](#) - Step-by-step, no technical background needed
- [The Marketer's Guide to OpenCLAW](#) - What it is, what MoltBot and CLAWDBot do, and why high-performance agents matter
- [Claude Marketplace, Dispatch, Code Channels](#) - Seven days of announcements explained

OFFICIAL DOCS

code.claude.com

Anthropic's official Claude Code documentation. Setup guides, feature reference, best practices, and the community skills database.

COMING SOON

The C10 AI Adoption Framework

A 12-chapter series covering every element of AI adoption. From Challenges (C1) through Compliance (C10), with implementation checklists and action plans.



Claude Code Implementation Guide

Practical AI implementation for business leaders

[AI Strategy](#)

[Agentic Automation](#)

[Digital Marketing](#)

[Training & Workshops](#)

anicca.co.uk • ann@anicca.co.uk • Leicester, UK • Est. 2007