finsi growth

✦ AI-Powered E-commerce Growth Agency ✦

Claude Code for Brands

THE AI OPERATOR FOR DTC GROWTH

How DTC operators put an AI agent to work inside Shopify, their ad accounts and their email — the stack it runs on, and how you build it for your brand.

WORKS INSIDE THE PLATFORMS YOU ALREADY RUN

Shopify Meta Google Klaviyo TikTok

Read this first

Most DTC founders have used AI. Almost none have put it to work.

You've pasted product descriptions out of a chat window. Maybe asked for subject lines, or a rewrite of an ad. That's AI as an advisor — you ask, it answers, and you carry the answer back into Shopify yourself. Useful. Capped.

Claude Code is a different thing. It's an agent. It runs on your computer, connects into the platforms your brand already runs on — Shopify, Meta, Google, Klaviyo, your analytics — and does the work. You describe the outcome in plain English. It carries it out, and nothing goes live without your approval.

At Finsi Growth we run client accounts through it every day: reports that build themselves, catalogs refreshed in brand voice, ad accounts that answer questions instead of hiding them. This guide is the map of that way of working, written for brand operators, not engineers.

It covers the WHY and the WHAT. Why copy-paste AI caps what you get. What the four parts of the stack are. How to work with an agent safely in live accounts. And a first hour you can run today, before you spend a dollar or change a single tool.

It doesn't cover the HOW for your specific brand — your catalog, your margins and your rules are yours. That's what a build-out is for, and the last page tells you how to get one.

One idea to hold onto as you read...

AI in a chat window is advice.
AI inside your stack is an operator.

Connect it to the work before you judge what it can do.

Copy-paste AI vs operator AI

Here's what we see across the brands we audit.

Most founders use AI like a faster freelancer. Draft this, summarize that, give me ideas. Each task gets 10 to 20% quicker, and the gain is real — but it stops there, because a human still ferries every output between the chat window and the tools. You become the API between your own systems.

A smaller group runs the business differently. They didn't ask for product descriptions — they gave the agent their catalog and a style guide, and now the catalog maintains itself in their voice. They didn't ask for a summary of last month — they rebuilt reporting, so Monday's trading report builds itself from live data. They didn't collect subject-line ideas — the welcome and win-back flows got drafted straight into Klaviyo for review.

Same AI. Same subscription price. Completely different outcome, because one group bought answers and the other installed an operator.

That's the shift this guide is about. Not doing the old jobs slightly faster — taking whole jobs off your desk, with you as the approver instead of the courier.

Copy-paste AI · a faster Tuesday	**Operator AI** · a different company
"Write 20 product descriptions, I'll paste them in."	The catalog refreshed in Shopify directly — in brand voice, from one approved style guide.
"Summarize last month from these screenshots."	Monday's trading report builds itself from live Shopify, ad and email data. Same format, every week.
"Give me some subject line ideas."	The welcome and win-back flows drafted in Klaviyo, ready for your review.

What Claude Code actually is

Most people think Claude is the same thing everywhere. It isn't.

There's the model — the intelligence — and there's Claude Code, an agent built on it. Same brain, completely different surface.

The chat version talks to you. Answers, drafts, research. It's a good advisor, one conversation at a time, and you do all the carrying.

Claude Code DOES the work with you. It runs on your computer, reads and writes real files, works inside your platforms through connectors, and can run several jobs in parallel. You brief it the way you'd brief a store manager: the goal, the context, what done looks like. It asks what it's missing, does the job, and shows you the result — in minutes, not days.

And no, you don't need to code. You type — or dictate — plain English. The "Code" in the name is about what it can touch, not what you have to write.

Chat AI · answers	**Claude Code** · outcomes
You copy answers back into your tools	**Works directly inside your accounts and files**
One thread of advice at a time	**Runs several jobs in parallel**
Forgets your brand between chats	**Briefed by a memory you control**
A clever intern	**An operator on your team — who always asks before acting**

One more thing matters before the next page: an agent is only as useful as what it can reach, and only as safe as the rules you give it. Both of those are things you set up — which is exactly what the stack is.

The stack: four parts

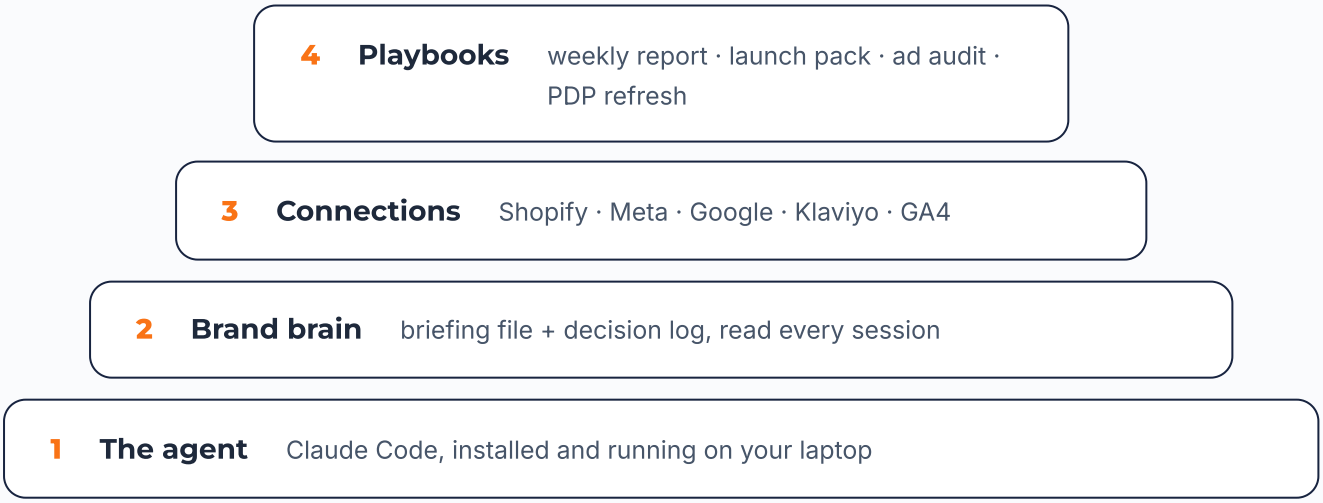
If you take one page from this guide, take this one.

Claude Code on its own is a capable agent with no memory of your brand, no connections and no repeatable way of working. That's not a flaw. It's the starting point.

What turns it into the way you run the brand is the stack around it. 4 parts.

1. **The agent** — Claude Code itself, installed and running on your laptop.
2. **The brand brain** — a briefing file it reads every session, and a decision log where every correction and rule gets saved for the next one.
3. **The connections** — the connectors that put it inside Shopify, the ad platforms, Klaviyo and your analytics.
4. **The playbooks** — jobs described once and run forever: the weekly report, the launch pack, the ad audit.

A playbook with no stack under it has nothing to run on. A reporting playbook can't report on accounts it isn't connected to. A copy playbook can't write in a voice it has no memory of. So you build the foundation first — then the playbooks make it pay, week after week.



Layers 1 to 3 are the machine. **Layer 4 is what your week actually feels like.**

PART 1

The agent

Part 1 is getting Claude Code running on your own computer.

It's one command, pasted into the terminal from the official quickstart at **code.claude.com**. Sign in with your Claude account — Pro or Max — and you're in. If a terminal feels like a leap, the Claude desktop app runs Claude Code too, and it's the easiest way to start.

From there, everything is plain English. Full sentences, the why, the context. Talking beats typing: dictation makes you explain, and explaining is what good briefs are made of.

Treat every session like briefing a capable new hire. Give it the job and what done looks like. It asks what it's missing. It does the work, saves the result, and walks you through it. The one difference from briefing a person is that it comes back in ten minutes, not next Thursday.

```
● ● ●   yourbrand / weekly  ·  claude

> Pull last week's sales vs the week before, flag any product whose conversion dropped, and
draft the Monday trading report in our format.
· Reading last 14 days of orders through the Shopify connector
· Checking ad spend against revenue by product through the ads connectors
· Running the weekly-report playbook with your brand style
✓ Saved reports/weekly-trading.html. 2 things to look at first, listed at the top.
```

Plain English in, finished work out. No code typed.

PART 2

The brand brain

Out of the box, the agent starts every session blank. Brilliant within a conversation, amnesiac between them. So you give it a memory — two layers, both yours to control.

The briefing file. A plain-text file it reads at the start of every session. Who the brand is. The voice, the bestsellers, the margins that matter. And the rules — above all, the one rule that makes everything else safe: **nothing publishes, sends, launches or spends without your explicit approval.**

The decision log. Where corrections and decisions get saved for the next session. "We never discount the hero product." "UK spelling, always." "That supplier issue is resolved — stop flagging it." Correct it once, and it's corrected forever on your brand.

That second layer is where the compounding starts. The 10th session is nothing like the 1st, because it arrives already briefed — your voice, your rules, your last five decisions, all in place before you type a word.

This is the part most people skip, and it's the part that decides whether the agent stays a toy or becomes the way your brand works. Playbooks before memory is the most common mistake we see.

LAYER 1

Briefing file

Who the brand is, the voice, the margins, the rules. Read at the start of every session.

LAYER 2

Decision log

Every correction and decision, saved for the next session. Nothing re-explained.

THE HABIT

Start + wrap

Load the memory at the start. Save what changed at the end. Every session, both ends.

Correct it once, and it's corrected forever.

Memory is what turns a clever tool into the way your brand works.

PART 3

The connections

Part 3 is putting the agent inside your platforms.

A connector is what lets Claude work directly in Shopify, Meta, Google Ads, Klaviyo, GA4 and the rest — reading real data, drafting real changes. Once a platform is connected, you can ask it anything in plain English. Pull last month's performance. Find the five products quietly losing conversion. Which campaigns spend without ever producing a second order?

This is the layer that turns a smart assistant into an operator — and it's where most setups stop short. The agent can only be as useful as what it can see.

If it touches revenue, connect it.

Here's a simple audit. For one day, write down every tab you open to answer a question about the business — Shopify admin, Ads Manager, Klaviyo, the spreadsheet of spreadsheets. That list is your setup to-do. Every line is a connection to make, and every connection you make removes that tab-switch forever.

Two kinds of access, and the difference matters. **Read access** answers questions and is safe by default. **Write access** ships real work — and in this stack, write always waits for your approval. The agent reads freely, and writes only when you say so.

Commerce + email

Shopify · Klaviyo · subscriptions · reviews

Paid media

Meta Ads · Google Ads · TikTok Ads

Analytics + files

GA4 · Search Console · your sheets and docs

Where a platform has no connector yet, exports on a schedule do the job. Same result: the agent sees the whole business.

PART 4

The playbooks

Part 4 is where the stack turns into your week getting lighter.

A playbook is a job described once and run forever. The Monday trading report. The product-launch pack — PDP copy, email flow, ad drafts, all from one brief. The monthly ad-account audit. The quarterly catalog refresh. Instead of re-explaining a 20-step process every time, you write it down once, and from then on one sentence runs the whole thing.

Playbooks grow. They start as a short instruction file — one task, one output, minutes to build. They pick up a template and data sources, so every run comes out structurally identical. The big ones become multi-step workflows that used to be a week of agency back-and-forth and now ship in an afternoon.

And here's the order that works, because it's the opposite of what most people do: **don't start by building the playbook.** Do the job in a live session first. Get the output to a standard you'd actually ship. Then say "turn what we just did into a playbook" — and run it on the next job of the same shape. The 2nd or 3rd time you do a job by hand is the moment to systemize it. Not before.

Every playbook you build is a problem you never solve by hand again. The library grows, and the 10th playbook takes a fraction of the effort of the 1st, because the stack under it is already there.



The 2nd or 3rd time you do a job by hand is the moment to systemize it. Not before.

What it solves

Here's what the stack takes off a brand operator's desk, day to day. Every one of these used to be an evening, a hire or a "next quarter." Now it's a session.

Product pages

BEFORE Rewriting descriptions one by one, someday

NOW **Catalog refreshed in brand voice from one style guide**

Weekly reporting

BEFORE An evening of screenshots into a doc

NOW **Builds itself from live data every Monday**

Ad account questions

BEFORE A PDF from the agency, once a quarter

NOW **Ask the account anything, any day, in plain English**

Email flows

BEFORE "On the list" for six months

NOW **Drafted in Klaviyo for review in one session**

Landing pages

BEFORE A brief, a designer, a developer, weeks

NOW **Drafted and previewed in a day, live when you approve**

Customer insight

BEFORE A thousand reviews nobody has read

NOW **Reviews mined into the angles your next campaign uses**

Competitor watch

BEFORE Occasional screenshots in a Slack channel

NOW **A standing brief, updated on a schedule**

Store health

BEFORE Stockouts and broken links found when they hurt

NOW **Daily checks that flag issues before they cost revenue**

Same recipe every time: a connection, a playbook, your data, the finished output.

Why it compounds

Here's why this changes the shape of your output, not just the speed.

Every session makes the next one smarter. You brief it, it works, and it saves what it learned. A correction made in week 1 is still holding in month 6. A decision made about your brand is there the next time you open it. Each session isn't a fresh start — it's a continuation, with the entire history behind it.

Now multiply it. Run three sessions at once and that's three jobs moving — the report, the launch copy, the ad audit — each drawing on the same memory. Your day stops being doing the work and becomes reviewing finished work and directing the next step.

And the playbook library compounds the same way. Every problem you solve by hand and turn into a playbook is a problem you never solve by hand again. The setup cost keeps dropping; the output keeps rising.

From the outside it looks like speed. From the inside it's a system that gets better every single day you use it — and never forgets.



Session 1	Session 5	Session 20	Session 50	Session 100+
Blank. You explain everything.	Knows the brand, the accounts, the rules.	Knows every decision and every correction.	The first draft arrives right.	The way you work. Nothing forgotten.

What the agent already knows about your brand, session after session. Every correction stays. Every decision carries forward.

Run it like we run client accounts

An agent inside live accounts deserves house rules. These are the ones every Finsi-managed account runs on, and they're the reason "AI in the stack" is calm instead of terrifying.

The one rule that does most of the work. One line in the briefing file: nothing publishes, launches, pauses, sends, discounts or spends without explicit confirmation. Read access is safe. Write access is real. The agent reads first, and writes with approval.

Approve each step until you trust it. For the first weeks, it asks before every action and you approve each one. That's how you learn what it actually does. Once you know its habits, you loosen the leash — on the low-stakes jobs first.

You are the quality layer. It builds fast, but it can't judge whether a recommendation suits your brand, and numbers deserve verification. Every deliverable gets a human review before it ships. After every launch, open the platform and check with your own eyes.

Push back. The first answer often isn't the best one, and "that can't be done" is usually wrong. Ask for another way. Ask what a senior operator would change. Be stubborn on the thing that matters.

Keep the keys out of the chat. Passwords and API keys never live in the conversation — connectors and a password manager hold them. Where a platform offers a connector, use it instead of a raw key.

Anything you run without thinking becomes a routine. The daily store-health check. The weekly report draft. The overnight data pull. If it runs on a clock and produces the same shape of output every time, it doesn't need you in the chair — and it still waits for approval before touching anything live.

Read access answers questions.
Write access waits for you.

That single line in the briefing file is what makes the rest of this guide safe.

Where to start: your first hour

You can run this today. No connections, no setup project, no migration. Claude Code, a folder, and a job you already owe someone.

1. **Install it.** One command from the official quickstart at code.claude.com, sign in with a Claude Pro or Max account, and you're in. Or start in the Claude desktop app — it runs Claude Code too, and it's the gentlest way in.

2. **Make a folder and open Claude inside it.** Drop in whatever you have — last month's Shopify export, your brand guidelines, the price list. First ask: "read every file in this folder and tell me what's here, one line each."

3. **Do one real job off your own to-do list.** Something that's been there a while and that you know what done looks like. The trading summary you owe your co-founder. Ten product descriptions in your voice. The competitor scan you've postponed for months. Say it with context: "Here's a job that's been on my list for weeks: [what it is, who it's for, what done looks like]. Ask me anything you need to know before you start, then do it."

4. **Review it like a contractor's first draft.** Talk to it the way you'd talk to the person who handed it to you: "shorter", "we never discount the hero product", "that number's wrong, it's 40". Then do the second one: "now turn that into an email to the team." The second and third asks are where it starts to feel different.

The rule for your first month: anything you'd have opened a doc or a spreadsheet for, do it in here instead.

After the first hour: connect a platform, set up the memory the first time it forgets something that costs you 20 minutes, and turn the first repeated job into a playbook. On your own, you'll get to the full stack in a few months. That path works. It's just slower.

The shortcut is having it built for you. As part of a growth audit, we map your brand's highest-value playbooks, set up the memory and the connections, and hand you a stack that's already working. Free for qualified brands.

Book your free growth audit → cal.com/finsi/audit

Takes 2 minutes to book. We'll map your brand's first three playbooks on the call.

finsi growth

Data-first growth for ambitious e-commerce operators

340% average ROAS across managed accounts	**+22%** LTV increase within the first 90 days
$4.2M client revenue generated in 2025	**−45%** churn through AI-driven retention

Book your free growth audit

Takes 2 minutes at **cal.com/finsi/audit** · We'll audit your stack, your cohorts and your growth leaks — and map the first playbooks worth building for your brand.