The AI Bootcamp

28

Build Your First AI Automation

In this chapter you'll learn

Trigger, action, result, and one small thing running by tomorrow morning

The AI Bootcamp

THE PROBLEM

Most AI education is a pile of disconnected tips. A LinkedIn carousel here, a YouTube tutorial there, a $2,000 cohort course that goes out of date before the refund window closes. None of it is built to be read in order, so nothing compounds. You finish twelve videos and can recite twelve facts, not perform one new skill.

THE FIX

The AI Bootcamp is built the other way. 90 skills, packed 2 to 3 per chapter, 36 chapters, 3 tracks. Every chapter states what it needs from you and what it hands to the chapter after it, so skill 47 only works because skill 12 already landed. Read it in order and by the end you are not collecting facts about AI, you are running work through it.

LETTER FROM THE EDITOR

We built this because we kept meeting smart people who felt behind on AI for no reason other than nobody had organized the material for them.

This book does not assume you are technical, and it does not water anything down either. Every chapter earns its place: a real skill, a real exercise, a one-page companion you can flip back to in month three when you have forgotten the details but still need the answer.

If a chapter loses you, or a chapter changes how you work, I want to hear about it either way.

Kris

Editor, AI Central

Write to me at **editor@thecentral.ai**

THE 3 TRACKS

AI 101

Get running

Pick a tool, set it up, write a clear instruction, use it safely at work.

AI 202

Prompt properly

The framework behind every good instruction, and catching AI when it is wrong.

AI 303

Build without watching it

Connect AI to your tools, and automations that run while you are not looking.

IN THIS LESSON YOU'LL LEARN

By the end of this chapter you will have one automation that has run at least once without you, you will be able to describe any automation in any tool as a trigger, an action and a result, and you will be able to tell the difference between one that had nothing to report and one that quietly died.

Everything in this track so far still needed you in the room. You connected it to your work in Chapter 26 and you decided what it was allowed to do in Chapter 27, and both of those happen with you sitting there watching. This is the chapter where something useful happens while you are asleep, and the first time that works it changes what you think this tool is for.

CONTENTS

28.1 Build the harmless one first, not the important one

28.2 Every automation is a trigger, an action and a result

28.3 One specific automation, running before you turn the page

28.4 Test it on purpose, and build it so silence means one thing

TOPICS REQUIRED

Chapters 25, 26 and 27. You need your assistant open, whatever you connected in Chapter 26, and the permission decisions you made in Chapter 27. You also need 15 uninterrupted minutes, because this is a build and not a read.

AFTER THIS CHAPTER

Chapter 29 assumes you have one automation that has already run without you, and that you know where to look to prove it ran.

28.1 Build the harmless one first, not the important one

Your instinct is to automate the thing that costs you the most time. That instinct is correct and it is the wrong first move, for a reason worth 2 minutes of your attention.

The one you want

The weekly report, the client update, the thing that eats 2 hours every Friday and that people notice when it is late.

TRAITS

- Worth real money in time saved
- Writes into a place other people read
- Somebody is expecting the output

PROS

+ The payoff is obvious and large
+ You already know exactly what it should say

CONS

– A misfire is visible to colleagues
– You learn the tool and carry the risk at the same time

The one to build now

A short daily brief on 3 things you would otherwise look up yourself, delivered to you and nobody else.

TRAITS

- Reads only, and writes nothing anywhere
- Arrives with exactly one person: you
- Useful enough that you will notice if it stops

PROS

+ A total failure costs you one deleted email
+ Same trigger, action and result as the big one

CONS

– Saves you minutes, not hours
– Feels too small to bother with, which is the point

Fig 28.1 Two automations, same machinery, different stakes. The one on the left is worth more and is why you are here. The one on the right is the one that teaches you the shape without charging you for the lesson.

28.1

The blast radius test. Ask what happens if it fires at 3am with completely wrong content. If the answer is "I delete an email", build it.

Reads only. A first automation that looks things up and reports back cannot damage anything it looked at.

One reader. It arrives with you. Making output land where colleagues look is Chapter 31, and that chapter wants a working automation first.

Here is what you are about to do, and it is worth naming before you do it. You are going to think of the single most annoying recurring job in your week, the one that eats a Friday afternoon, and you are going to try to automate that first. Everybody does. It is the right list to be thinking about and the wrong place to start.

The problem is not difficulty. It is that you would be learning what these things feel like when they work, and what they look like when they fail, at the same moment as trusting one with something that matters to other people. Those are two separate jobs and doing them together is how people end up with a broken automation they are afraid to touch and too embarrassed to mention.

So your first one has a different purpose. Its job is not to save you time. Its job is to teach you the shape, prove to you that the thing actually runs when you are not looking, and show you what a failure looks like from the outside. It should be small enough that you would build it again from scratch without complaining.

Use the blast radius test to choose it. Imagine it fires at the worst possible moment with completely wrong content. If the honest answer is that you delete an email and shrug, you have a candidate. If the answer involves apologising to a client, that one is not first.

In practice that means 3 properties. It reads things rather than changing them, so nothing it touches can be damaged. It writes nowhere: no sending, no posting, no calendar entries, no rows added to anything. And it arrives with exactly one person, you, so a mistake never becomes somebody else's problem.

The valuable one is not lost, it is Chapter 29, which is entirely about finding the highest-value recurring work in your week and automating that. It is the same machinery with the stakes turned back up. For now, think of 3 things you would genuinely like to be told about every morning, and keep reading.

28.2 Every automation is a trigger, an action and a result

Three parts, in every tool, under a dozen different names. Learn the shape here and the next platform change costs you 10 minutes instead of a weekend.

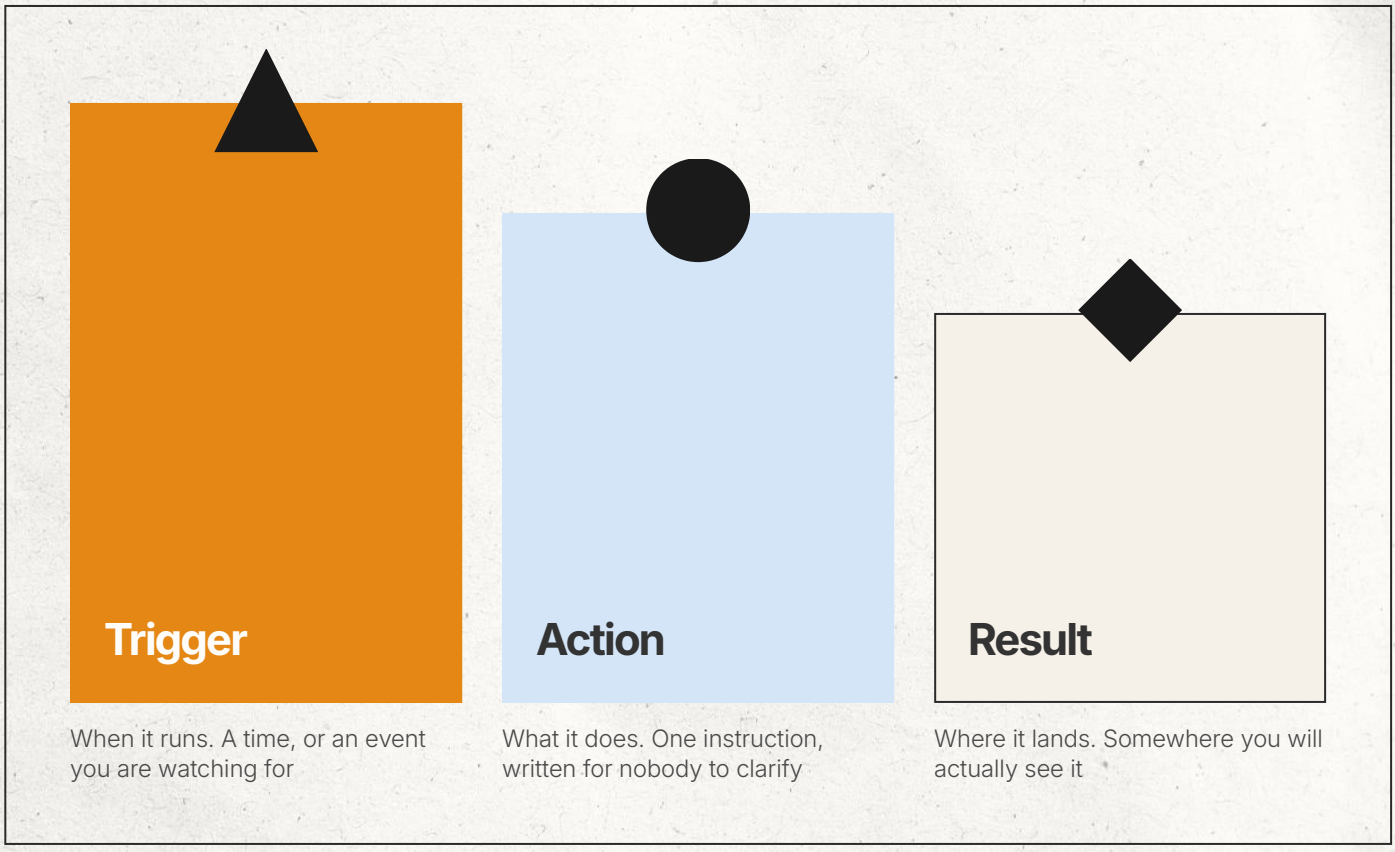


Fig 28.2 The 3 parts, in the order they happen. Every automation you will ever build is these 3 and nothing else. The bar heights are visual rhythm, not importance: leave any one of them vague and the whole thing fails.

28.2

Trigger. A clock, or an event: a message arriving, a file appearing, a form submitted. Start with a clock, because a clock is testable.

Action. The instruction, written the way Chapter 5 taught you. Nobody is there to answer a follow-up question, so vagueness costs more here.

Result. The output has to arrive somewhere you already open. A result that lands in a folder you never visit is the same as no result.

Every tool in this category names these things differently, and the naming is the only genuinely confusing part. Zapier calls the whole thing a Zap. Make calls it a scenario. ChatGPT calls it a scheduled task, Microsoft calls it a scheduled prompt, Google calls it a scheduled action, and Power Automate calls it a flow. Underneath the vocabulary, all of them are 3 parts in the same order.

The trigger is what starts it. Triggers come in two kinds: a time, or an event. A time trigger fires at 7am whether or not anything interesting happened. An event trigger fires when something specific occurs, such as an email arriving from a named sender. Your first automation should use a time, because a clock is predictable, which means you can test it, and because an event trigger that never fires is indistinguishable from one that is broken.

The action is the work, and it is a prompt like any other. Everything Chapter 5 taught you about context, task and format still applies, with one difference that matters more than it sounds: there is nobody there to ask a clarifying question. If a human would have needed to check something with you before starting, the automation will simply guess, in silence, every single day.

So write the boundary into the instruction itself. As well as saying what it should do, say what it must not do: do not email anyone, do not post anything, do not add anything to my calendar, do not change any file. This feels excessive on day one. It is the single habit that separates people who run 12 automations calmly from people who ran 3 and turned them all off.

The result is where the output lands, and it is where most first automations quietly fail. Not because the work was wrong, but because the answer arrived somewhere the person never opens. Decide this deliberately: a push notification, an email to yourself, a page in the tool you check every morning with your coffee. If you have to remember to go and look, you will not.

Before you open anything, write your 3 lines on paper. When it runs. What it does. Where it lands. If you cannot write those 3 lines, you do not yet have an automation, you have an intention.

28.3 One specific automation, running before you turn the page

A morning brief on 3 things you would otherwise look up by hand. It is genuinely useful, it reads only, and a total failure costs you nothing.

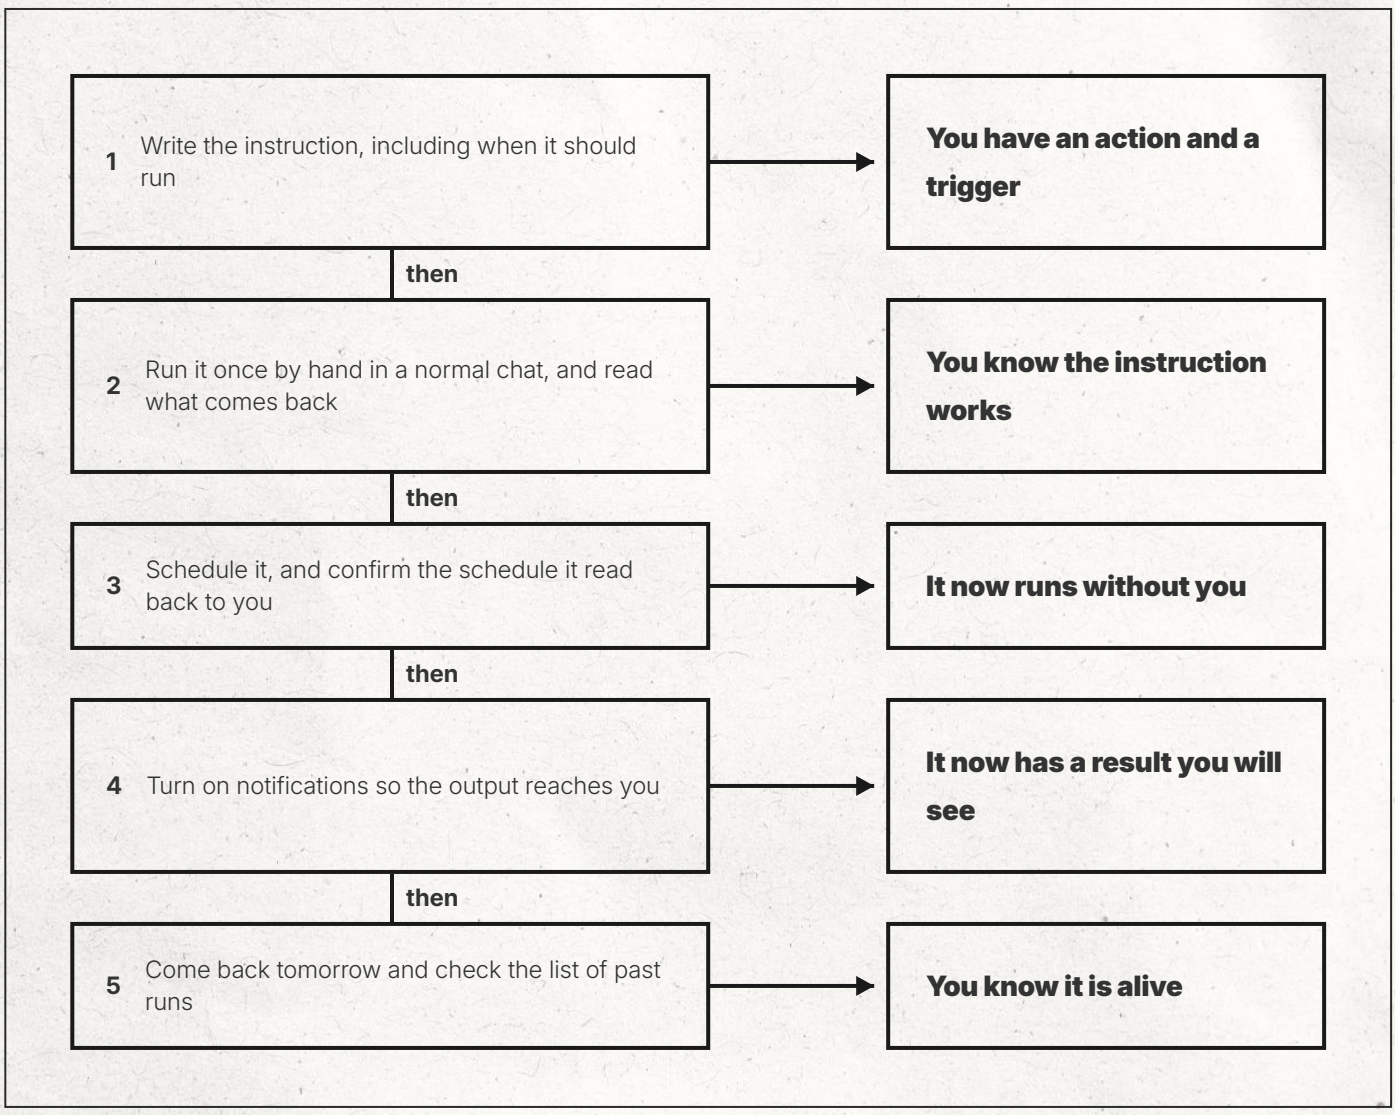


Fig 28.3 The build, in order. Step 2 is the one everybody skips, and it is the one that catches almost every mistake while the mistake is still cheap.

28.3

ChatGPT A page called Scheduled. Included on the free plan: up to 3 active tasks, one run a day, in a morning, afternoon or night window. Paid plans add hourly runs and exact times.

Claude Scheduled tasks live in Cowork, on paid plans only. Hourly, daily, weekly, weekdays or on demand. They run on Anthropic servers, so they run with your laptop shut.

Gemini Scheduled actions need a paid Google AI plan or a qualifying Workspace plan. Spark, the version that works continuously in the background, is top tier only.

Copilot Scheduled prompts come with a work Copilot licence and run in Teams, Outlook and Office.com chat. You manage them from the Automations tab.

Here is the one to build. Pick 3 things you currently check by hand and would rather be told about: a competitor, a large client, a regulator, a category you are supposed to have an opinion on at the next board meeting. Name them specifically. Vague subjects produce vague briefs.

Now the instruction, near enough word for word. "Every morning, search the web for anything new in the last 24 hours about these 3 things: [name], [name] and [name]. Give me at most 5 bullets. One sentence each, with a link. If there is genuinely nothing new, reply with the words nothing new and stop. Do not email anyone, do not post anything, do not add anything to my calendar."

Read that again and find the 3 parts. Every morning is the trigger. Search and summarise into 5 bullets is the action. The reply that arrives on your phone is the result. The last sentence is the boundary, and the sentence before it is doing something cleverer than it looks, which is the subject of the next section.

Where you type it depends on which assistant you picked in Chapter 1, and the notes above give you the specifics. The one thing worth saying out loud: a free ChatGPT account can do this today, with up to 3 of these running at once, one run a day, delivered in a rough morning window rather than at an exact minute. If that is what you have, you need buy nothing to finish this chapter. Claude, Gemini and Copilot all put scheduling behind a paid or work plan.

If your assistant gates it and you are not paying, you have a free route that teaches the same anatomy possibly better. Zapier's free plan costs nothing, allows unlimited automations, and limits each one to 2 steps and 100 runs a month. Two steps is exactly one trigger plus one action, which is the shape you just learned, and a once-a-day automation fits inside 100 comfortably. You describe what you want in a sentence and it drafts the automation for you to check.

Go and build it now, before you read the last section. It takes about 4 minutes, and the last section makes considerably more sense once something of yours is scheduled.

28.4 Test it on purpose, and build it so silence means one thing

An automation you cannot check is not an automation, it is a rumour. Two small habits turn it into something you can actually rely on.

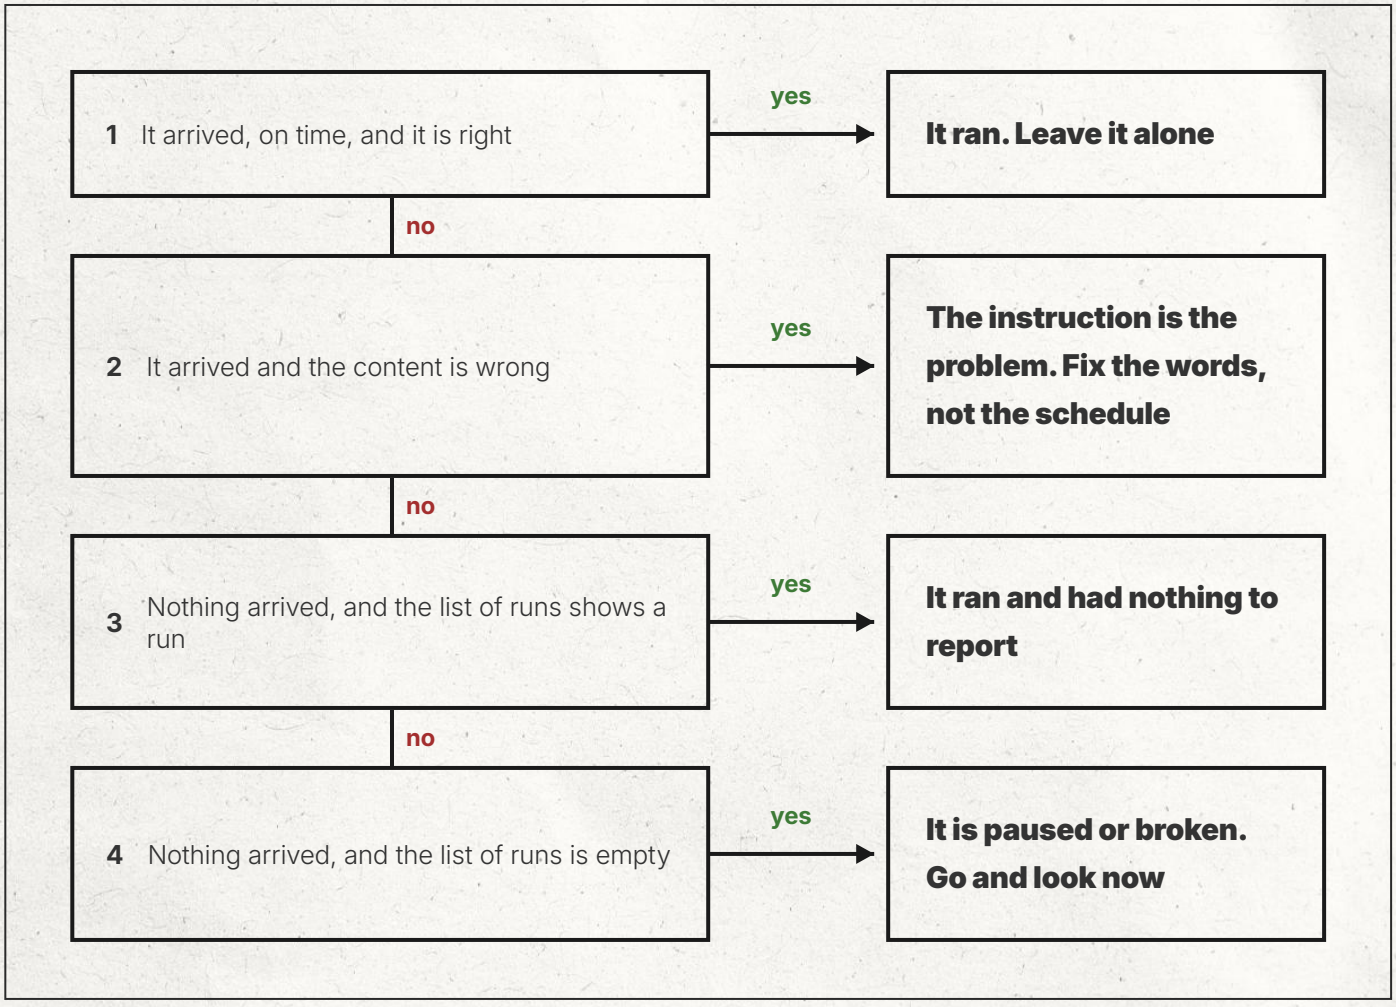


Fig 28.4 Read down until a line is true. From your inbox the bottom two rows look identical, which is exactly why the instruction has to say something even when it has nothing to say.

28.4

Run it by hand first. If the instruction does not work while you are watching it, it will not work while you are asleep.

Make silence mean one thing. Tell it to say nothing new when there is nothing new. Then no message ever means no run.

Check the run list, not your memory. Every one of these tools keeps a list of past runs. That list is the evidence. Your impression is not.

The first habit is deliberate testing, and it is 30 seconds of work. Before you schedule anything, paste the same instruction into an ordinary chat and read what comes back. You are not checking whether the answer is interesting. You are checking whether the instruction produces the shape you asked for: 5 bullets, one sentence each, with links, and nothing else attempted. If it does not do that with you watching, scheduling it only means being wrong on a timer.

The second habit is the one almost nobody is taught, and it is why that instruction ends the way it does. An automation that reports nothing when there is nothing to report is indistinguishable from an automation that died 3 weeks ago. Both look like an empty inbox. You will assume the quiet one is working, right up until the day you needed it and it had not run since the spring.

So make silence mean one thing. Tell it to reply with the words nothing new rather than staying quiet, and now the rule is simple: something arrives every day, and a day with no message is a day to go and look. That single clause converts an invisible failure into an obvious one, and it transfers to every automation you will ever build, in any tool, forever.

Then know where the evidence lives. Each of these tools keeps a list of past and upcoming runs: ChatGPT on its Scheduled page, Claude on the Scheduled page in Cowork's sidebar, Copilot in the Automations tab. Look at that list tomorrow rather than relying on whether you remember seeing a notification. Notifications get swiped away half asleep.

It is also worth knowing what quietly stops these things, because it is rarely dramatic. ChatGPT pauses a task if it goes unused for long enough, if it needs a decision from you, or, and this one catches people, if you delete the chat the task was created in. A revoked connection, a changed password or an expired sign-in will stop any of them. None of that announces itself.

Two things you should deliberately not do yet. Do not let this one take an action out in the world: sending, posting and filing come later and they come with rules. Supervising an automation that does act is Chapter 34, and what to do when one breaks overnight is Chapter 35. Stringing several apps together is Chapter 33. Right now you want one small honest thing running, and tomorrow morning you want to check that it did.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 20 minutes and one real decision that opens Chapter 29.

Do these before Chapter 29

3 exercises. About 20 minutes

Before you open anything, write your automation out in 3 parts. One line each, in plain words. If a box will not fill, that is the part that is not decided yet.

TRIGGER: WHEN IT RUNS	ACTION: WHAT IT DOES, AND WHAT IT MUST NOT DO	RESULT: WHERE IT LANDS, AND WHO ELSE CAN SEE IT

Now go and build it. Run the instruction once by hand in a normal chat before you schedule it, then schedule it. Write down the exact sentence you used, so you can rebuild it in a minute if it ever disappears.

Tomorrow, after it has run without you, write down 2 things: what actually arrived, and where you looked to confirm it ran rather than assuming it did.

Chapter 28 on one page

THE 3 PARTS	Trigger, when it runs · Action, what it does · Result, where it lands. Every automation in every tool is these 3, whatever the buttons are called
A SAFE FIRST ONE	Reads only · writes nowhere · arrives with one person, you. If it misfires at 3am you delete an email and shrug
MAKE SILENCE MEAN ONE THING	Tell it to reply "nothing new" when there is nothing new. Then a day with no message always means a broken automation, never a quiet week
WHERE SCHEDULING LIVES	ChatGPT: Scheduled, included free, once a day · Claude: Scheduled, inside Cowork, paid plans · Gemini: scheduled actions, paid Google AI plans · Copilot: scheduled prompts, with a work licence · Or Zapier free, 2 steps and 100 runs a month
NEEDED BEFORE CHAPTER 29	One automation that has run without you at least once, and the run list you checked to prove it

BRING TO CHAPTER 29

The automation you built, the exact sentence you gave it, and the first thing it sent you. Chapter 29 starts from something that already works and asks the harder question of what deserves the next one.

BOTTOM LINE

You have one automation running without you, you can name its trigger, its action and its result, and you can tell whether it ran.

NEXT CHAPTER

5 Questions to Find Work Worth Automating

Chapter 29 turns the stakes back up. It walks you through an audit of your own recurring work, so you can see which repeated jobs are actually worth automating, and then builds the one that pays for the rest.

BRING WITH YOU

- ☐ Your working automation from this chapter, and proof that it ran
- ☐ A rough list of the jobs you repeat every week, however untidy

GO FURTHER

The Ultimate AI Library

Every guide, prompt pack and playbook AI Central has published, in one place. The reference shelf this course keeps pointing you to.

explore the library ↗



WHO MAKES THIS

AI CENTRAL

AI Central Media is a London-based editorial team. Our flagship publication covers practical AI for senior professionals, and 100+ AI companies, SaaS platforms and growth teams have advertised with us since 2023.

100+

AI companies, SaaS platforms, education brands and growth teams have advertised with us since 2023

613K

Accounts reached every month, across the 7 channels we operate

151

Countries our readers are in, plus all 50 US states

PROUD PARTNER OF

GAMMA · IIElevenLabs · guidde. · Outskill · Luma AI

HubSpot · attio · Synthflow · Notion · Replit

We turn attention into pipeline.

visit thecentral.ai ↗