The AI Bootcamp

26

Connect AI to the Tools You Already Use

In this chapter you'll learn

What a connector actually buys you, where the plan wall sits, and the one system to connect first

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 A LinkedIn carousel here, a YouTube tutorial there, a $2,000 cohort course that goes out of date before the refund window closes.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 can say exactly what a connector gives you that attaching a file does not, find out in 2 minutes whether your own plan even offers one, name the single risk that comes with standing access and the settings that contain it, and you will have connected exactly one system on purpose.

You have spent 25 chapters carrying material to the assistant by hand. Every conversation starts with you going to find the file, and finding the file is most of the work. This chapter is about the assistant going to look instead, which sounds like a small convenience and is not: it changes which questions you bother asking at all.

CONTENTS

26.1 Attaching a file is a loan. Connecting a system is a key

26.2 All four connect to something. Only two connect to their own house

26.3 The instruction can arrive inside the thing it reads

26.4 Connect the system you already waste the most time searching

TOPICS REQUIRED

Chapters 7, 9 and 25. You need the habit of putting your own material in front of your assistant, the answer you got in Chapter 9 about what your organisation actually permits, and Chapter 25's distinction between an assistant that replies and one that does something.

AFTER THIS CHAPTER

Chapter 27 assumes at least one system is connected, because permission to act only becomes a real question once there is something to act on.

26.1 Attaching a file is a loan. Connecting a system is a key

Chapter 7 had you hand over one document for one conversation. This is the other thing entirely: standing access to a whole system, so the assistant goes and finds the document itself.

Attaching a file

You find the document, you upload it, and the arrangement ends when the conversation does.

TRAITS

· One file, chosen by you, for one conversation
· Nothing persists after you close the tab
· Works on every plan of every assistant

PROS

+ You know precisely what it saw
+ No standing access to anything

CONS

− You do the searching, every single time
− You can only give it what you already thought to give it

Connecting a system

You authorise once, and from then on it can search that system on your behalf until you revoke it.

TRAITS

· A whole system, not one file
· Live: it sees what is there today, not a copy
· Persists until you disconnect it

PROS

+ It finds material you had forgotten existed
+ Questions you would not have bothered asking become cheap

CONS

− Standing access, granted once, easy to forget about
− It reads material written by other people, which matters, see 26.3

Fig 26.1 The same material, reached two different ways. What changes is not what it can read, it is who does the looking and how long the arrangement lasts.

26.1

Chapter 7's move. One file, chosen by you, for one conversation. It ends when you close the tab.

This chapter's move. Standing access to a whole system, live, until you go and revoke it.

Your permissions, not new ones. A connector reaches what you can already reach. If you cannot open a file, neither can it.

Here is the problem this chapter exists to fix. By now the assistant is genuinely useful, but every session starts the same way: you go and find the deck, the contract, the last three emails on the subject, and you paste or attach them. The retrieval is the work. The assistant only starts once you have finished doing the hard part.

A connector removes that step. You authorise it once against a system you already use, your file storage, your mail, your calendar, your notes, and from then on it can go and look. Ask it what your team agreed with a supplier in March and it searches, finds the thread, and answers from what is actually there rather than from what you remembered to attach.

The most important fact about connectors is also the least well known, and it should settle most of the nervousness people feel at this point. A connector inherits your own permissions. Anthropic's documentation states it about as plainly as it can be stated: if someone cannot access a specific file, channel or record in the source system, the connector cannot reach it either. Connecting does not create access. It borrows the access you already had.

What you genuinely buy, then, is not a smarter assistant. It is the disappearance of the retrieval step, and the real effect of that is on the questions you ask. Nobody digs through 4 years of shared drives to settle a small question. When the digging costs nothing, small questions get asked, and small questions are where most of the wasted time in a senior job actually sits.

What it costs is standing access. An attachment is over when the conversation is. A connection sits there quietly until you remember to remove it, and the thing on the other end of it keeps changing without you. That is not a reason to avoid connectors. It is a reason to grant a narrow one deliberately rather than 5 broad ones casually, which is what section 26.3 is about.

Before you read on, go and open your assistant's settings and find the list of connectors or apps. Do not connect anything yet. Just look at what is on offer, because what is on offer turns out to depend heavily on what you are paying, and that is the next section.

26.2 All four connect to something. Only two connect to their own house

The categories on offer are the same everywhere: file storage, mail, calendar, chat, and notes or project tools. What differs sharply is which plan you need, and whose building you are standing in.

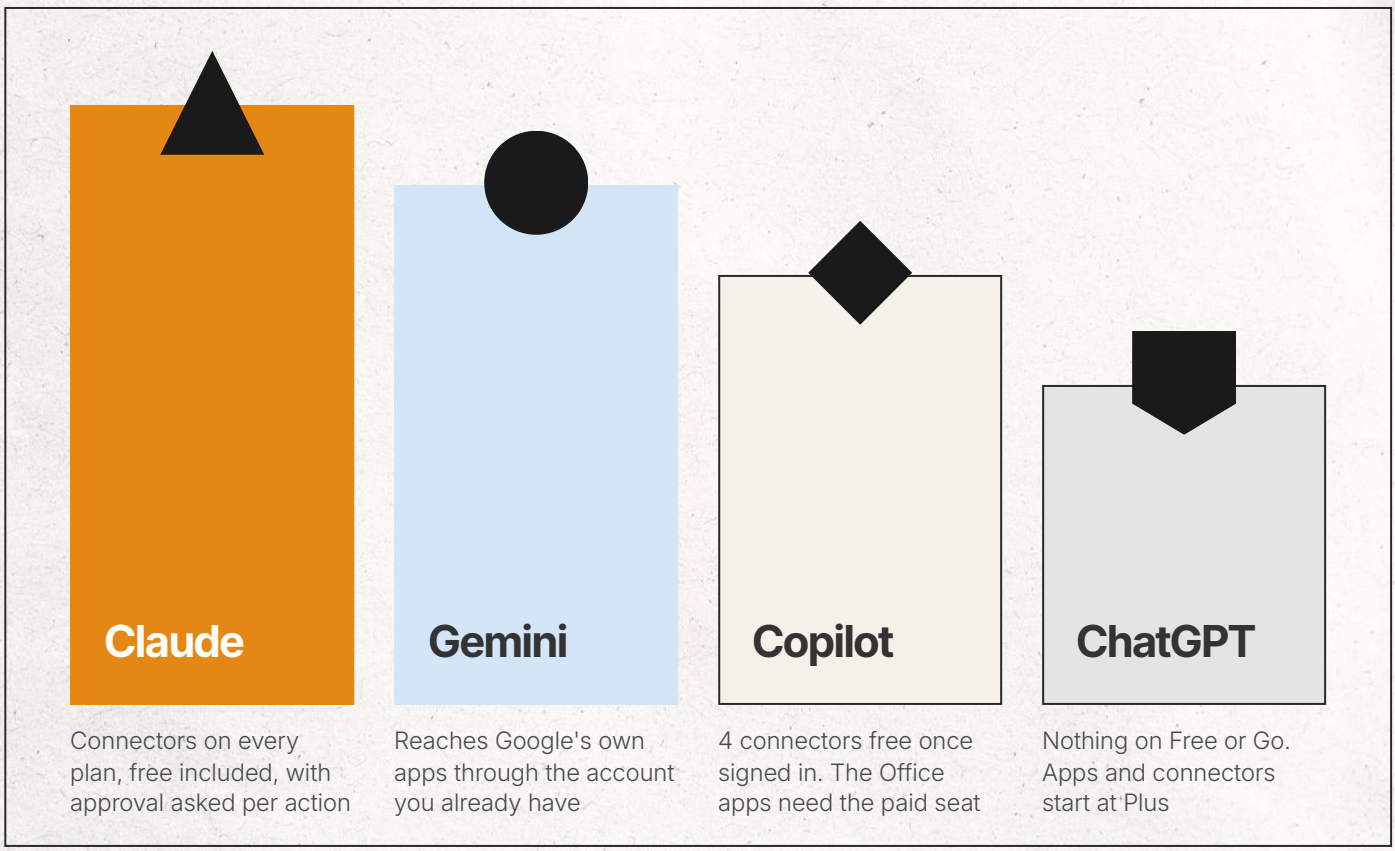

Fig 26.2 How much you can connect on an ordinary personal account, before anyone buys you a work licence. The heights encode that and only that: they are not a ranking of the assistants, which is Chapter 1 material.

26.2

ChatGPT Filed under apps. OpenAI's own plan comparison lists connecting to your other tools as unavailable on Free and on Go, and available from Plus upwards.

Claude Google Workspace and Microsoft 365 connectors on every plan, free included. On Team and Enterprise an owner has to switch one on before anyone can use it.

Gemini Connected Apps on a personal Google account reach Gmail, Drive, Docs, Calendar, Keep, Tasks and more. Gemini inside Gmail and Docs is a separate, paid Google AI plan.

Copilot Free Copilot connects OneDrive, Outlook.com, Google Drive and Gmail. Word, Excel, Teams and company SharePoint need the paid Microsoft 365 Copilot seat.

Every one of the four connects to roughly the same 5 categories, and if you read only the marketing you would think they were interchangeable. They are not, and the difference is almost entirely about money and about whose software you are already inside.

Claude is the most open here and it is worth saying so plainly, because it is the opposite of what people assume about the paid tiers. Its Google Workspace connectors, Gmail, Calendar and Drive, and its Microsoft 365 connector are available on every plan including the free one. It asks your approval before each action by default. If you are on a company Team or Enterprise account the shape changes: an owner has to enable a connector for the organisation before you can authorise it yourself, which is a policy decision, not a technical one.

ChatGPT is the opposite case and the one most likely to catch you out. Its connectors are now filed under apps, and OpenAI's own plan comparison lists apps that connect to your other tools as not included on the free plan and not included on Go. They begin at Plus. So if you have been reading about ChatGPT reaching into Google Drive and cannot find the setting, nothing is broken: you are on the wrong plan. Business and Enterprise accounts add something consumers cannot get at all, an administrator syncing an entire shared drive or SharePoint site so the whole company searches the same indexed material.

Gemini and Copilot have a structural advantage inside their own ecosystems, and it is fair to state it plainly rather than pretend the four are level. Gemini does not really connect to Gmail and Drive so much as already be in the same account as them, which is why the Google side of Gemini works on an ordinary personal account with nothing purchased. Copilot has the same relationship to Microsoft 365. The two advantages are not equally available to you, though, and that is the twist worth knowing: Gemini's is yours to switch on this afternoon, while Copilot's real reach, into Word, Excel, Outlook, Teams and your company's SharePoint, sits behind a per-seat licence your employer buys. Free Copilot connects 4 things and stops. What your organisation is willing to license is Chapter 9's territory, and it is a conversation with IT rather than a setting.

One more thing is happening underneath all of this and it is worth 30 seconds even though it is technical, because it explains why the list of connectable things keeps growing without anyone announcing anything. There is now a common standard for wiring an assistant to an outside system, called the Model Context Protocol. Anthropic published it and gave it away, and its own documentation compares it to USB-C: one plug shape, so a tool that builds the socket once can be reached by any assistant that speaks it. It is not one vendor's project any more. Claude, ChatGPT and Gemini all reach systems through it, along with most of the developer tools, and the specification is revised on a public schedule.

You do not need to understand a word of how it works. The consequence is the part that matters to you: when your project tool or your CRM announces support for it, that is the day it becomes connectable to whichever assistant you chose, without either company having built anything specifically for the other. So go and look at your own connector list again with the plan wall in mind, and find out which of the systems you actually use are on it.

26.3 The instruction can arrive inside the thing it reads

Standing access has a real cost, and it is not the one most people brace for. The risk is not that a company reads your mail. It is that a document you did not write can talk to your assistant.

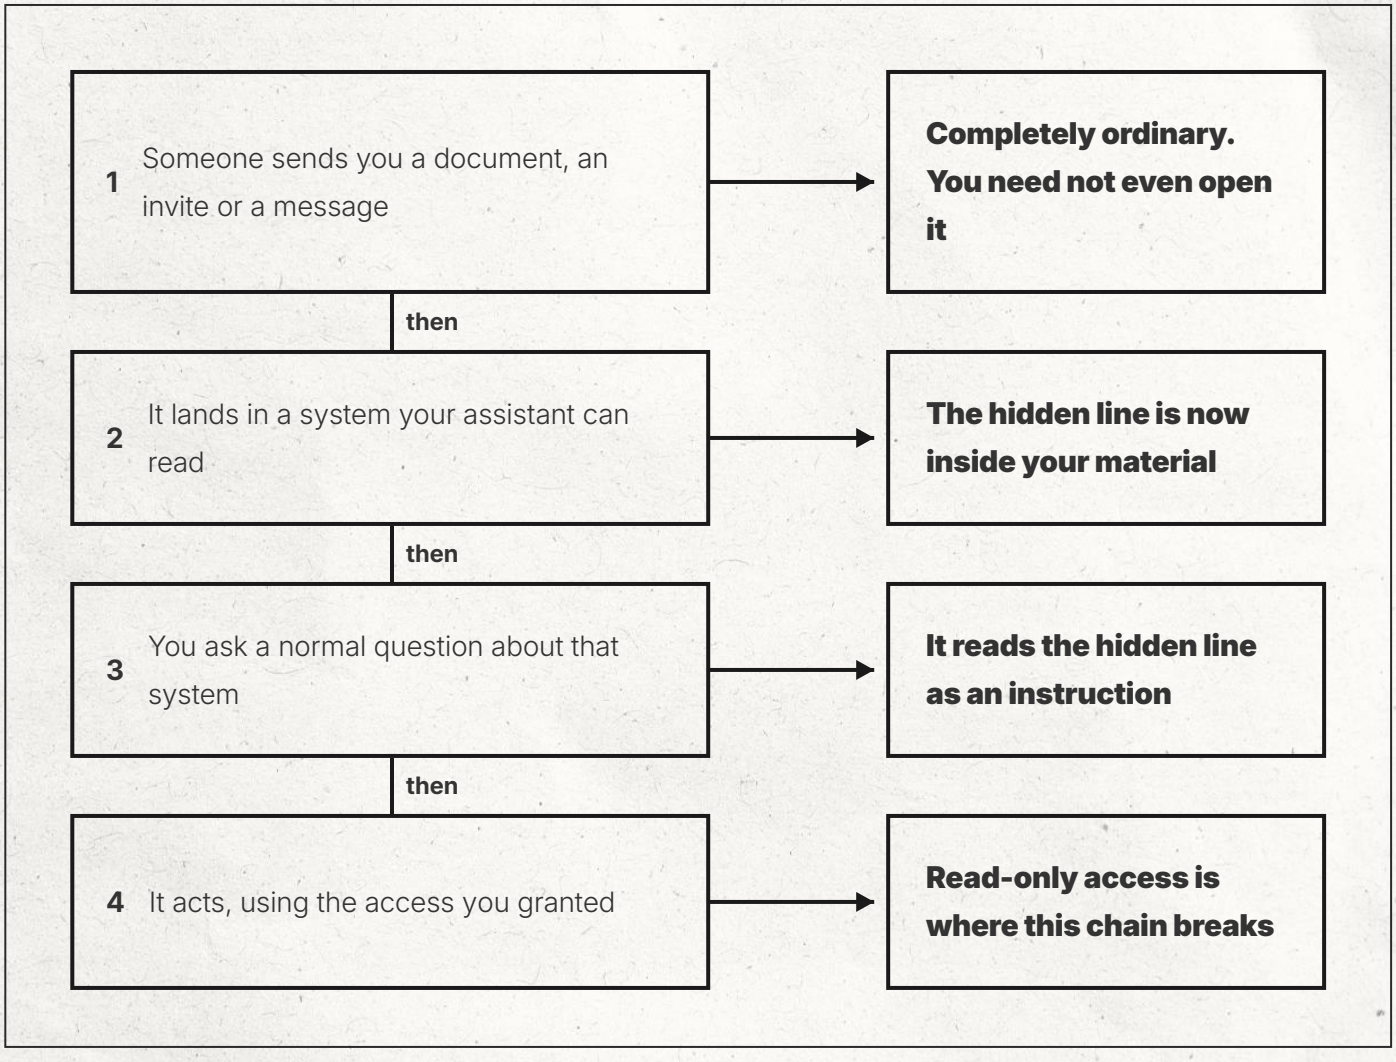

Fig 26.3 How an ordinary-looking message becomes an action. Notice step 3: you did nothing unusual. This is why read-only access is worth so much more than it sounds.

26.3

The mechanism. It cannot reliably tell your instruction from an instruction written inside something it was asked to read.

Where it stands. OWASP says no fool-proof prevention is known. OpenAI calls it a frontier research problem. Treat it as a property, not a bug awaiting a patch.

What you control. One system, read-only or ask-before-acting, narrow questions, and revoke anything you stopped using.

Now the honest part, and it deserves to be neither skipped nor dramatised. Connecting an assistant to your mail or your files grants real access, and there is one specific failure mode you should understand before you grant it.

The assistant reads everything as words, and it cannot reliably tell the difference between an instruction you typed and an instruction sitting inside a document it was asked to read. Someone who knows this can write a line into a calendar invite, a shared file or an email, phrased as a command and often invisible to you, and hope that your assistant reads it while doing something ordinary on your behalf. The industry calls this indirect prompt injection. The plain version is that the content can give orders too.

Two documented cases show the shape of it, and both are more reassuring than they first sound. In 2025 researchers found a way to pull information out of Microsoft 365 Copilot with no click from the victim at all. It was rated critical, given a CVE number, and Microsoft's own advisory records that it was fully fixed on their side before disclosure, with no action required from anyone using the service. Separately, a team from Tel Aviv University, the Technion and SafeBreach hid instructions in Google Calendar invitations, so that asking Gemini to summarise the week set them running. They reported it privately in February 2025 and presented it that August, by which time Google had shipped detection for it and started requiring confirmation on sensitive actions.

Read those two the right way. In both cases researchers found the problem, disclosed it, and the vendor fixed it, which is the system working. What you should not take from them is that the category is closed. OWASP's own guidance says it is unclear whether any fool-proof prevention exists, and OpenAI describes the problem as a frontier research challenge rather than a solved one. So the sensible posture is neither alarm nor indifference: assume this is a permanent property of connecting an assistant to material other people can write into, and give it less to work with.

That means granting the narrowest access that does the job, and every one of the four gives you the controls to do it. Connect 1 system, not 5. Prefer reading to acting: ChatGPT's permission ladder runs from always ask, through allowing reads while still asking before any change, up to allowing everything, and leaving it at one of the first two costs you almost nothing. Claude asks for approval before each action by default, and on a company account an owner can permit reading while blocking writing outright. Then ask narrowly. OpenAI's own advice is specifically to avoid instructions shaped like review my emails and take whatever action is needed, because an open-ended instruction is exactly what an injected one can hide inside.

Two last practical habits. Revoke what you stopped using: disconnecting lives in the same settings panel as connecting, and a connector you forgot about is pure exposure with no benefit. And keep separate the question of whether you are allowed to connect a work system at all, which is not a security question but a policy one, and which Chapter 9 already answered for your organisation. What it means to let it act rather than read is Chapter 27, next.

26.4 Connect the system you already waste the most time searching

There is a right first connection and it is rarely the impressive one. It is wherever your own retrieval is worst, which for almost every senior professional is not the inbox.

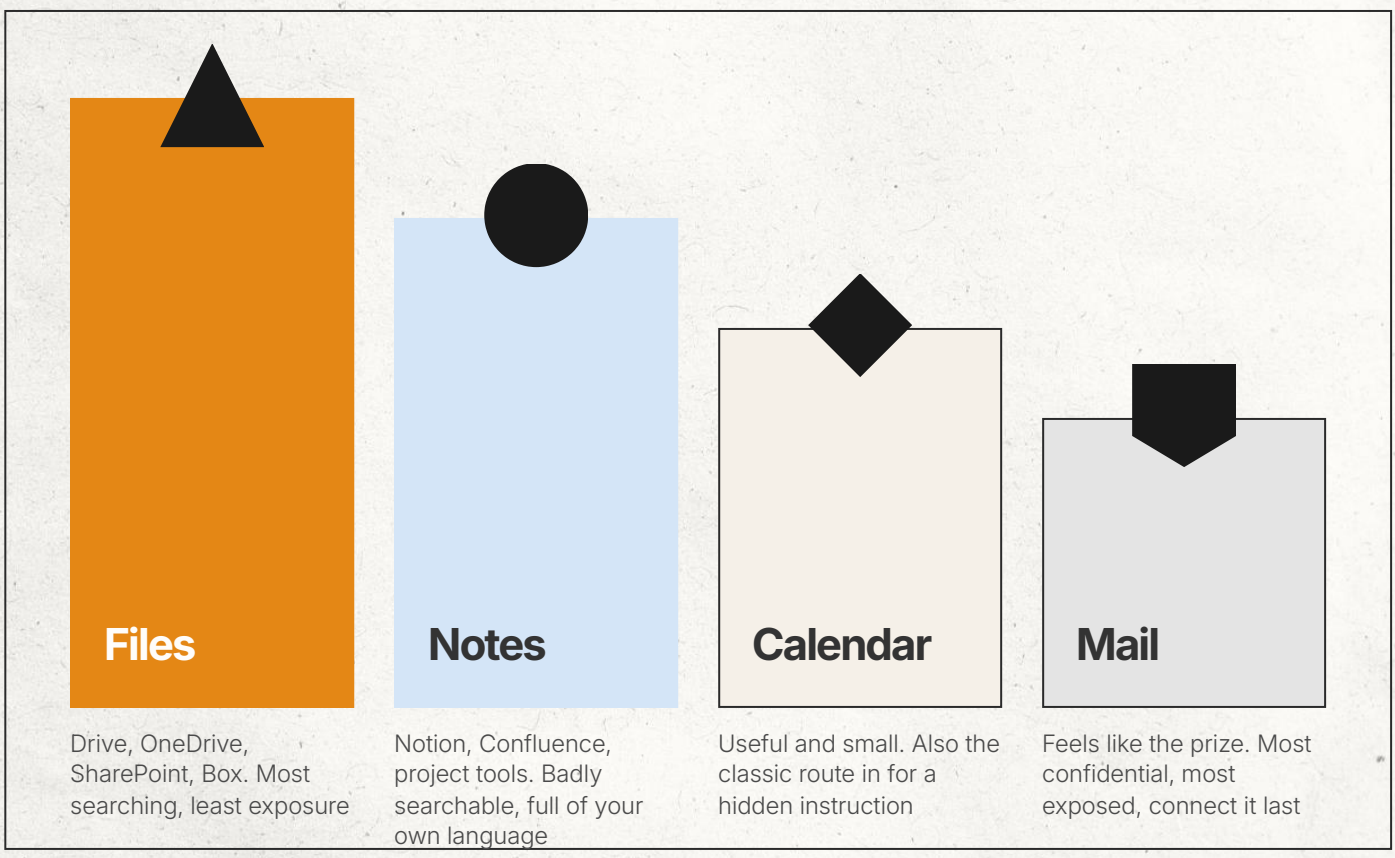

Fig 26.4 The order to connect things in, weighing what you get against what you expose. The heights encode that recommended order, not how useful each system is in your job.

26.4

Start with storage. The most time recovered, the least exposure, and almost everything you want from it is reading.

Leave mail until last. Most confidential material, most text written by strangers, and a write goes out with your name on it.

One at a time. Live with a connector for 2 weeks. If you never once asked it about that system, disconnect it.

So which one first. The answer is not the one that demonstrates best in a meeting. It is the system where you currently waste the most time looking for things, and if you are honest about your own week that is almost never your mail.

File storage is the first connection for most senior professionals, and it wins on both sides of the trade. On the payoff side, it is where your retrieval is genuinely bad: the deck from the offsite, the version of the contract that was actually signed, the analysis someone did in March that you can picture and cannot find. On the risk side, almost everything you want from it is reading, and nothing you ask it to do there leaves the building.

Second is whatever holds your written thinking: your notes tool, your wiki, your project tracker. The payoff is high for the same reason it is hard to search, because it is written in your organisation's own vocabulary rather than in English, and search boxes are bad at that. The exposure is a step up, though, because that material is largely written by other people, which is precisely the surface described in the previous section. Connect it read-only and you have taken most of the risk out of it.

Calendar comes third. It is genuinely useful and the payoff is modest, because you are not usually failing to find a meeting. It also carries two specific costs: an invitation is text a stranger can put in front of your assistant, which is exactly the route the Gemini research used, and write access there is visible to other people the moment it is exercised.

Mail is last, and this is the recommendation people argue with, so here is the reasoning. Your inbox has the highest concentration of confidential material you hold, it is the surface most exposed to text written by people you have never met, and write access means something can be sent under your name. It also has the weakest case on payoff, whatever it feels like. You rarely fail to find an email. You fail to find the document that was attached to it. The inbox feels like the prize because it is where you feel busy, and feeling busy is not the same as searching badly.

Go and connect exactly 1 thing, set it to read-only or to asking first, and then use it for a fortnight before adding another. Most people connect 5 systems in an afternoon and never ask any of them a question, which is all of the exposure and none of the benefit. What it means to let a connected system act rather than just answer is Chapter 27. Putting your inbox and calendar to work properly is Chapter 30, and getting output back into the documents your colleagues open is Chapter 31. All three assume you did this part first.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 27.

Do these before Chapter 27

3 exercises. About 15 minutes

Before you connect anything, write down the last 3 times you went looking for something and it took longer than it should have. Name the system each one was in, not the file.

WHERE I WAS SEARCHING	WHERE I WAS SEARCHING	WHERE I WAS SEARCHING

Open your assistant's settings and find the connectors or apps list. Write down which of the systems you just named is actually on offer to you. If none of them is, write down what would have to change: a different plan, or a decision by your employer.

Now connect exactly one, and set its permission to read-only or to asking before it acts. Then ask it one real question you would not have bothered to look up by hand.

WHAT I CONNECTED	THE PERMISSION LEVEL I SET	THE FIRST REAL QUESTION I ASKED IT

Chapter 26 on one page

ATTACH VS CONNECT	An attachment is one file for one conversation and ends with it. A connector is standing access to a whole system, live, until you revoke it
IT BORROWS YOUR ACCESS	A connector reaches only what you could already reach. It never creates permission you did not have
WHERE THE PLAN WALL SITS	Claude: every plan, free included · Gemini: your personal Google account already reaches Google apps · Copilot: 4 connectors free, the Office apps need the paid seat · ChatGPT: not on Free or Go, starts at Plus
THE ONE RISK WORTH KNOWING	Instructions can hide inside content it reads, so it acts on them with your access. Contain it with read-only, 1 system at a time, narrow questions, and revoking what you stopped using
NEEDED BEFORE CHAPTER 27	One system connected on purpose, its permission level set deliberately, and one real question already asked of it

BRING TO CHAPTER 27

The one system you connected and the permission level you chose for it. Chapter 27 starts from exactly that pairing, because permission to act only means something once there is something real to act on.

BOTTOM LINE

You can say what standing access buys and what it costs, find the plan wall on your own account, and you have connected one system deliberately rather than five casually.

NEXT CHAPTER

What AI Can Do Without Asking You

Chapter 27 takes the connection you just made and asks the harder question: what changes when the assistant stops reading and starts doing. Tool use, browsing, and what you are actually agreeing to when you grant permission to act.

BRING WITH YOU

- ☐ The system you connected in this chapter, and the permission level you set on it
- ☐ One task you would genuinely be willing to let it carry out rather than describe

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 ↗