AIWhy Briefing

AI Assistants

Define your AIWhy before you pick a tool

Briefing as of November 18, 2026

Questions, or want to talk it through? Talk to an Outcome Architect:
https://aiwhy.io/talk-to-an-outcome-architect

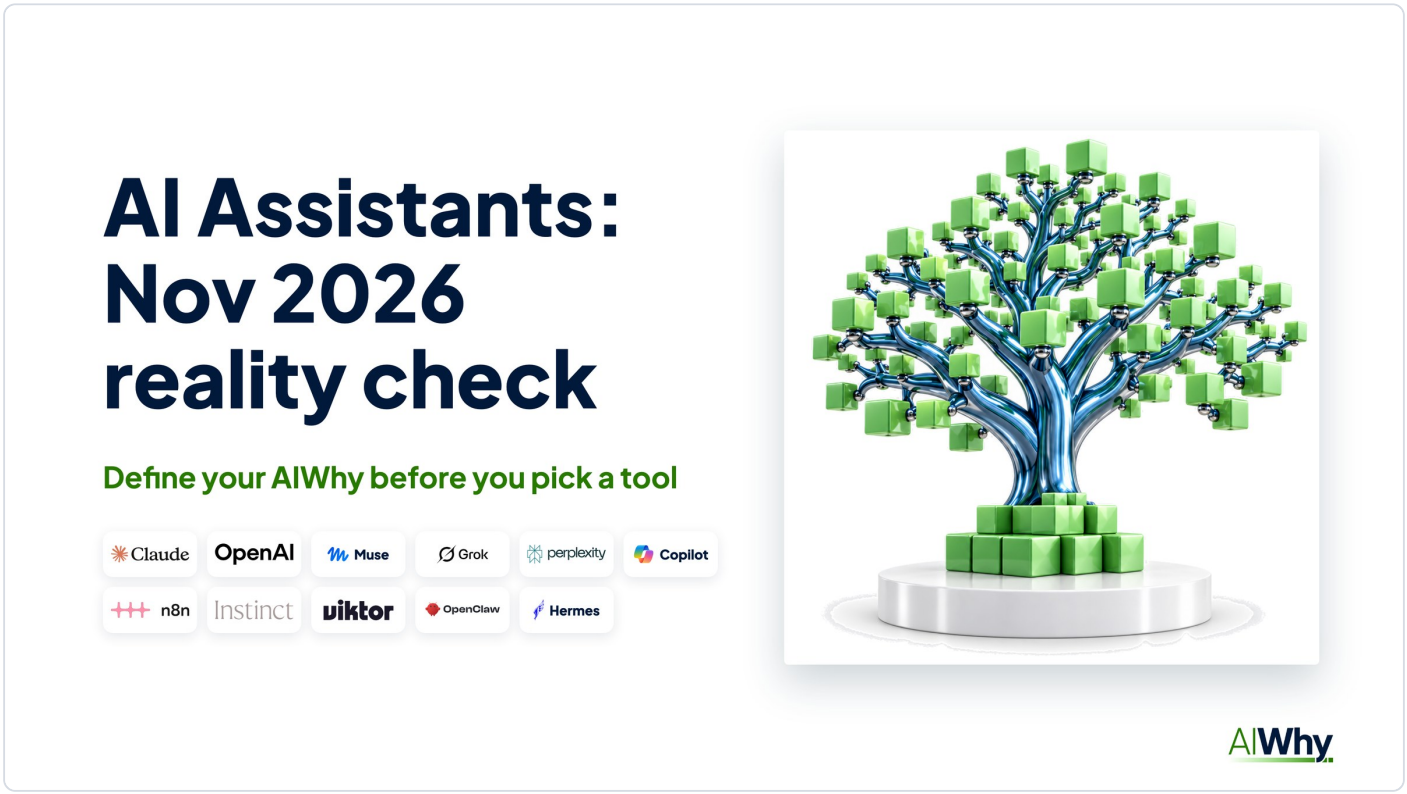


AIWhy

© 2026 Lextech. All rights reserved.

Contents

AI Assistants: Nov 2026 reality check 3
Capability Equity™ 4
"Which AI assistant should I use?" is the wrong question 5
Anatomy of an AI assistant 6
From a category of one to crowded 7
How much assistant comes built in? 8
What's actually working 9
Alex's dot, Rook 10
One owner, four Grok bots 11
What Tom learned 12
The pattern 13
Head to head: points out of 25 14
Vendor lock-in is a trap 15
The whiteboard and the filing cabinet 16
Write it down as a skill 17
Scheduled skill or AI assistant? 18
Paying does not buy privacy 19
Red alert 20
Interview your next AI assistant 21
Interview Prompt 22
We've learned this once already 24
Best practices 25
Key takeaways 26
Define the help you want. Don't accept the default. 27



AI Assistants: Nov 2026 reality check

This is AIWhy's look at always-on AI assistants as of mid-November 2026.

An always-on AI assistant is software that works for you in the background. It keeps going even when you are not talking to it. Think of it as a new employee who works on your behalf around the clock.

In the ten weeks before this briefing, Anthropic, OpenAI, Meta, xAI, Perplexity and n8n all released one. Microsoft and Google are close behind.

Everyone wants to know which one to use. Our advice is to decide why you want one first, then pick a tool. We call that reason your AIWhy.

These products are only weeks old, and their terms change every month. Check any number here against the maker's own website before you rely on it.

SOURCES
Introducing Claude Duet: https://claude.com/resources/articles/claude-duet
Introducing dots: https://openai.com/index/introducing-dots/
Introducing Muse: The World's First Personal AI Agent Built for Everyone: https://about.fb.com/news/2026/09/introducing-muse-personal-ai-agent/
SpaceXSI's Grok Bot 1.0 Adds Slack, Teams and Meeting Support: https://www.tokenpost.com/news/technology/55124
Perplexity launches Teammates: https://runtimewire.com/article/perplexity-teammates-computer-agents-ongoing-work
Meet Nod: https://blog.n8n.io/meet-nod



Capability Equity™

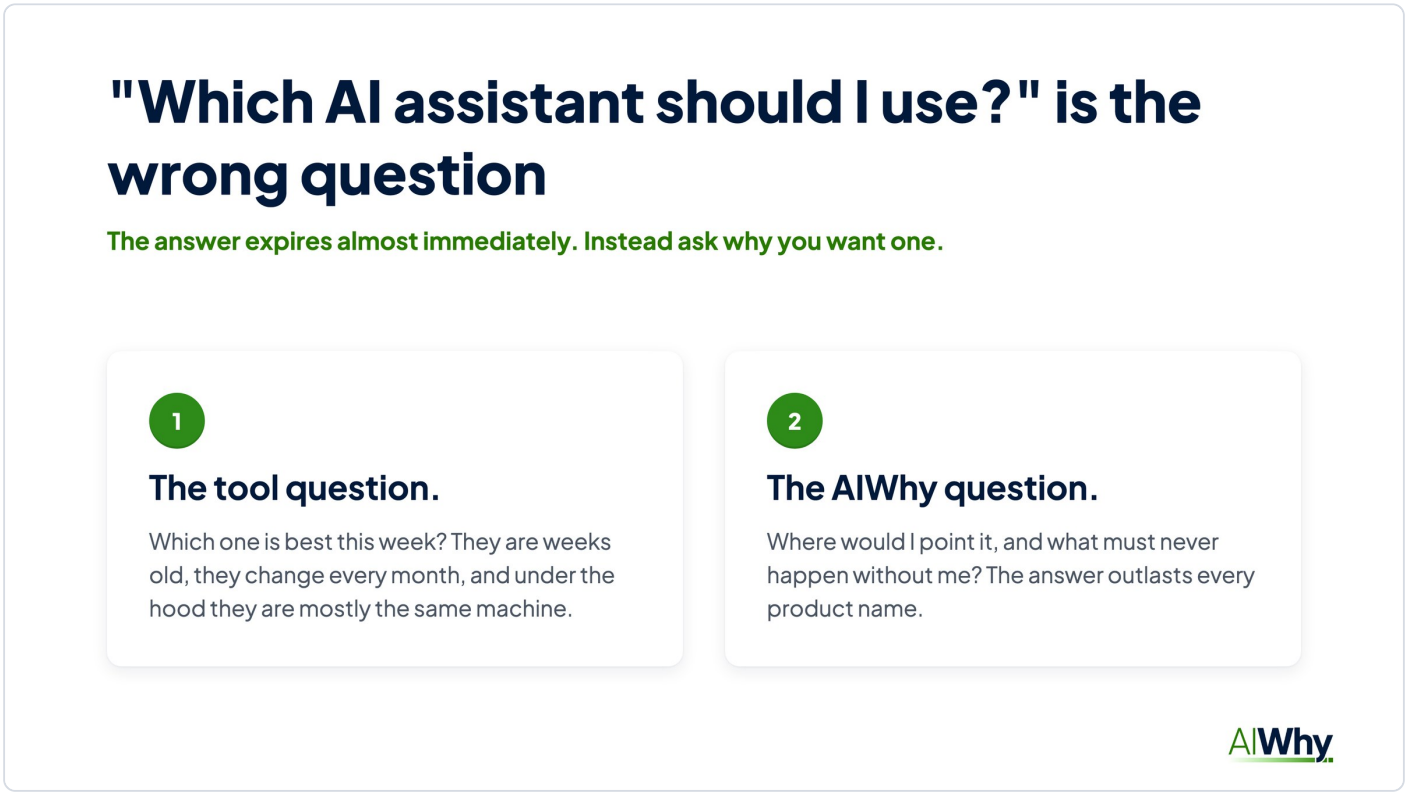
AIWhy works from a method called Capability Equity. Its main idea is to stop renting your competitive advantage.

Your know-how is what your business does better than others. It should be an asset the business owns. Today it lives in people's heads and in software you pay for each month.

The method rests on five ideas:

- Treat your know-how as an asset.
- Make it run. Runnable IP™ is your knowledge, written down so it can do the work.
- Use AI as your software. You use it and adjust it so it fits how you really work, instead of working the way a program forces you to.
- Build value the business owns. It can be checked, passed on, and it keeps growing.
- Grow with the team you have. Take on more clients and bigger projects with the same people.

For AI assistants, this gives a simple test. One is worth it when what you teach it becomes something your business owns.



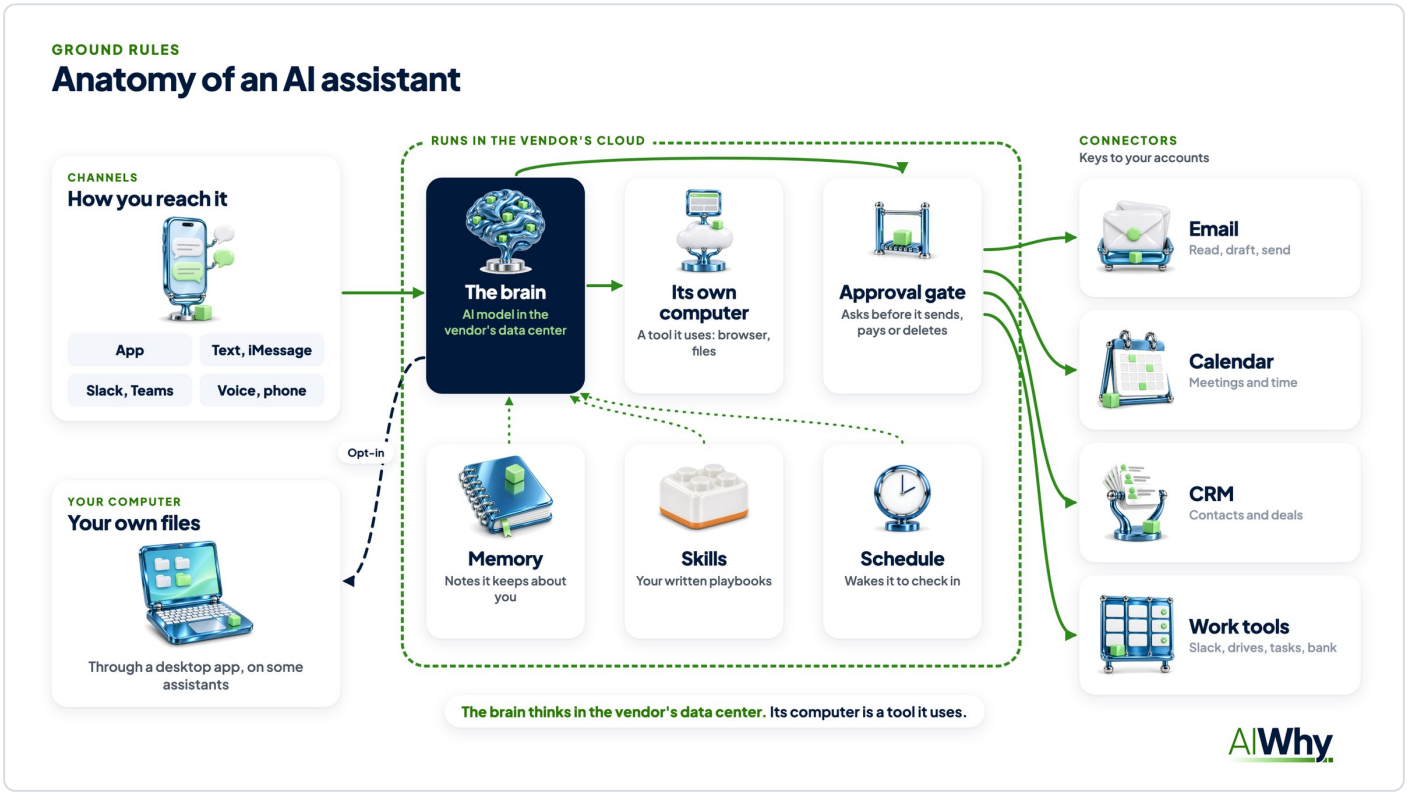
"Which AI assistant should I use?" is the wrong question

Most people ask which AI assistant is best. Any answer to that goes out of date almost at once. These products are weeks old and change every month. Underneath, they mostly run on the same kind of engine.

A better question is why you want one at all. What job would you give it? What must it never do without you? Those answers stay true no matter which product you pick.

Every success story in this briefing starts from one clear job: a morning summary, a search for new sales leads, a weekly check of timesheets. Nobody in the public record succeeded by asking it to "run my life."

The rest of this briefing follows three threads: where to point an assistant, what it must never do without you, and how to choose one.



Anatomy of an AI assistant

An AI assistant has nine parts. We use these names throughout.

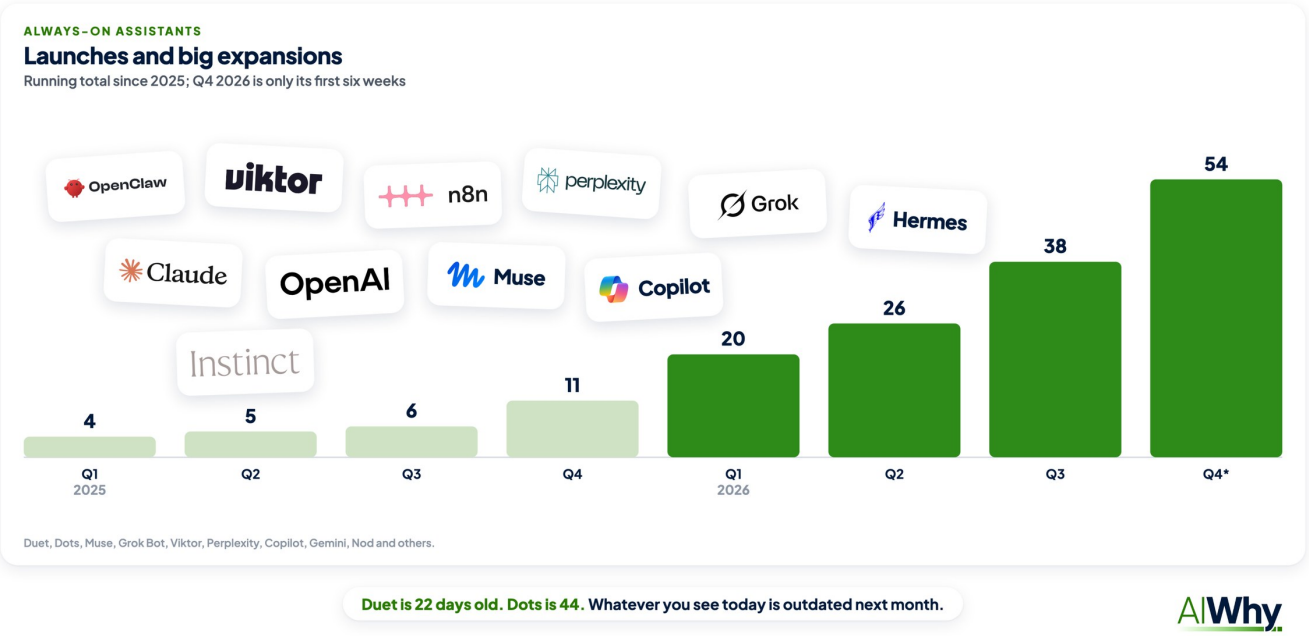
- **Channels:** how you talk to it. An app, a text, Slack or Teams, or your voice.
- **Brain:** the AI model, the part that reads, writes and decides. It is rented from an AI company.
- **Its own computer:** a virtual computer (one made only of software) with a browser and files. It uses this to click through websites as you.
- **Memory:** the notes it keeps about you.
- **Skills:** written step-by-step instructions for a job, like a playbook for a new hire.
- **Schedule:** an alarm clock that wakes it to check your inbox and calendar, so it can act before you ask.
- **Approval gate:** a stop to check with you before it sends, pays or deletes. On most products a second AI decides when to ask.
- **Connectors:** the keys that let it into your email, calendar and other business software.
- **Your computer:** four of the assistants can reach files on your own machine through an app.

SOURCES

Claude Duet Safety, Privacy, and Security: https://claude.com/resources/articles/claude-duet-safety-privacy-and-security
Dots privacy, security, and safety FAQs: https://help.openai.com/en/articles/20001529-dots-privacy-security-and-safety-faqs
Grok Bot for teams and enterprises: https://docs.x.ai/grok-bot/teams-and-enterprises
When Nod asks before acting: https://nod.n8n.io/docs/trust/approvals/

THE FIELD IS EXPLODING

From a category of one to crowded



Duet is 22 days old. Dots is 44. Whatever you see today is outdated next month.

AIWhy

From a category of one to crowded

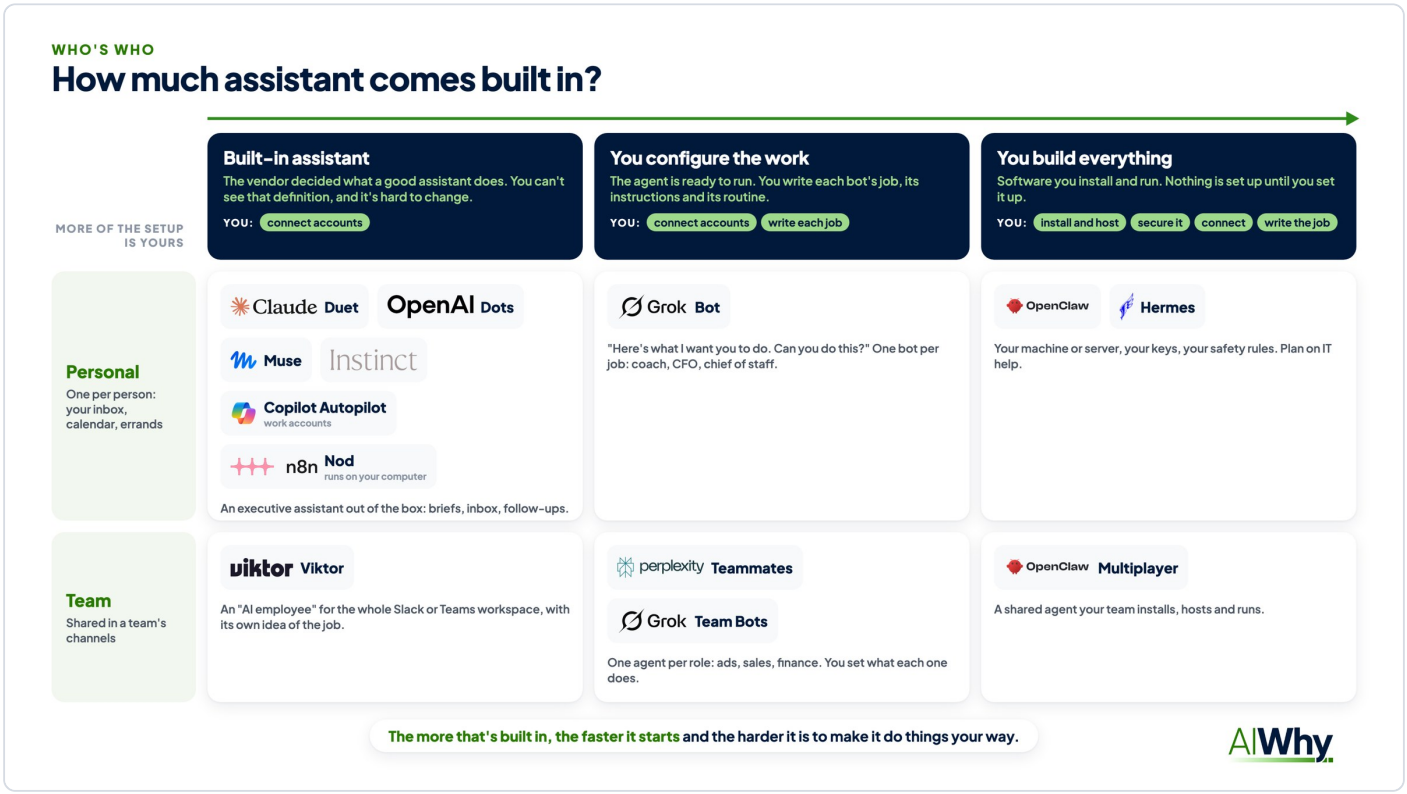
In all of 2025 there were eleven launches and big expansions of always-on assistants. By mid-November 2026 the total reached fifty-four. More than half came after July. Sixteen arrived in the first six weeks of the fourth quarter alone. That is more than in any earlier full quarter.

The best-known products are only weeks old. As of November 12, 2026, n8n's Nod was two days old, Perplexity Teammates a week, Claude Duet 22 days, OpenAI Dots 44 and Meta Muse 65. They're younger than most of the leftovers in your refrigerator. The oldest ones, Viktor and OpenClaw, aren't a year old yet.

November brought more. Dots turns on by default for large business customers on November 18. Microsoft's Autopilot is in early access, which means a first group of customers can try it.

SOURCES

Managing dots in ChatGPT workspaces: https://help.openai.com/en/articles/20001554-managing-dots-in-chatgpt-workspaces
What's New in Microsoft Copilot, October 2026: https://techcommunity.microsoft.com/blog/microsoft-copilot-blog/what%E2%80%99s-new-in-microsoft-copilot--october-2026/4568528
Introducing Claude Duet: https://claude.com/resources/articles/claude-duet
Perplexity launches Teammates: https://runtimewire.com/article/perplexity-teammates-computer-agents-ongoing-work
Meet Nod: https://blog.n8n.io/meet-nod



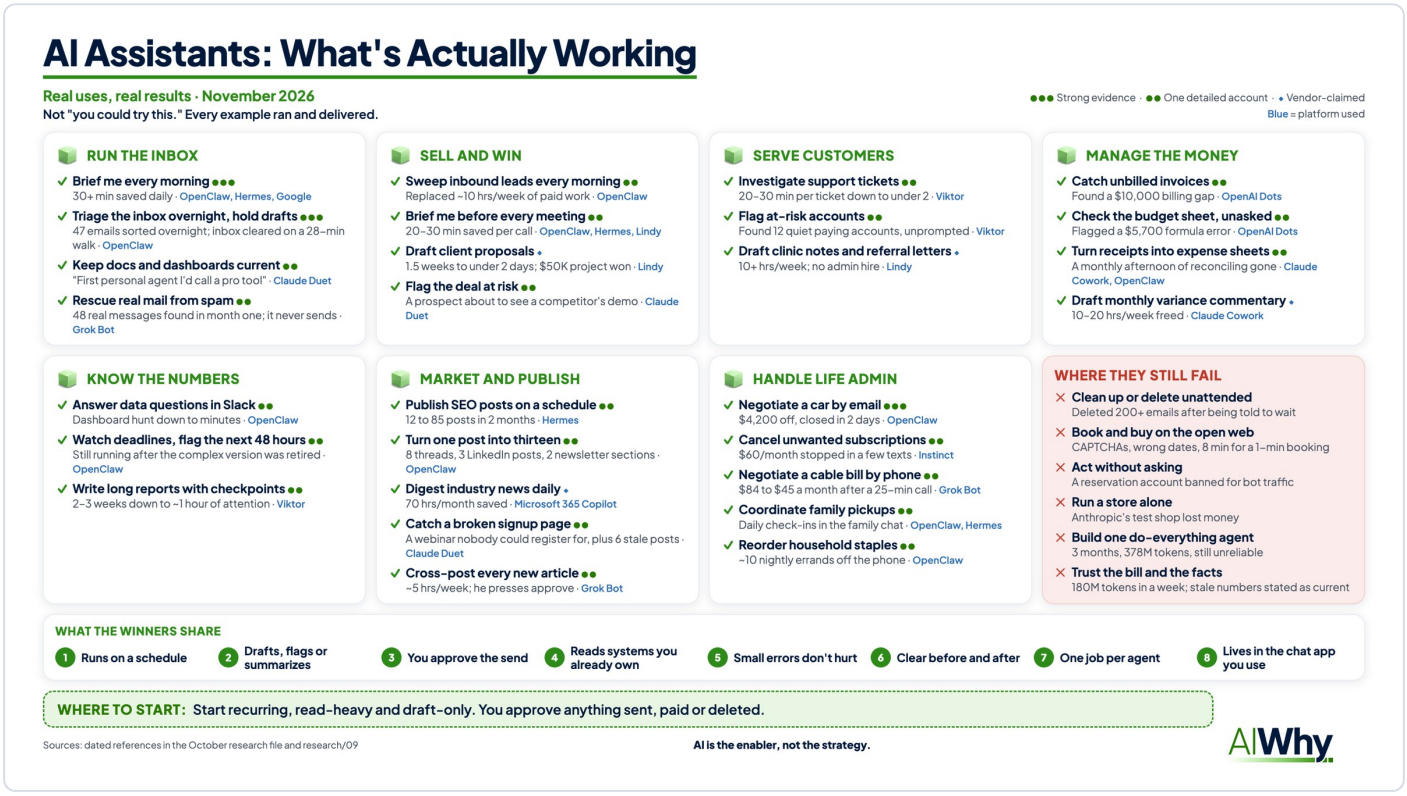
How much assistant comes built in?

These products differ in two ways. Some serve one person and some serve a team. Some come ready to go, and with others you do the setup yourself. The more that comes built in, the faster you start, and the less you can do things your way.

- **Built in:** Claude Duet, OpenAI Dots, Meta Muse and others. The maker decides what it does and what it watches. You connect your accounts and go. You can't see those choices or easily change them.
- **You set the job:** Grok Bot and Perplexity Teammates are ready to run, but you write each one's job yourself. You say where the work comes from, what it delivers and when it acts. Tom Schwab built four Grok bots this way: a coach, a finance chief, a chief of staff and a marketing chief.
- **You build everything:** OpenClaw and Hermes are programs you install and keep running yourself. You keep them secure and connect them to every account. Plan on IT help.

SOURCES

Get started with Perplexity teammates: https://www.perplexity.ai/help-center/en/articles/2026102601-perplexity-teammates
Teammates product page: https://www.perplexity.ai/hub/products/teammates
Grok Bot: Bots (memory, Team Bots, deletion): https://docs.x.ai/grok-bot/bots
OpenClaw GitHub README: https://github.com/openclaw/openclaw
Hermes Agent GitHub README: https://github.com/NousResearch/hermes-agent



What's actually working

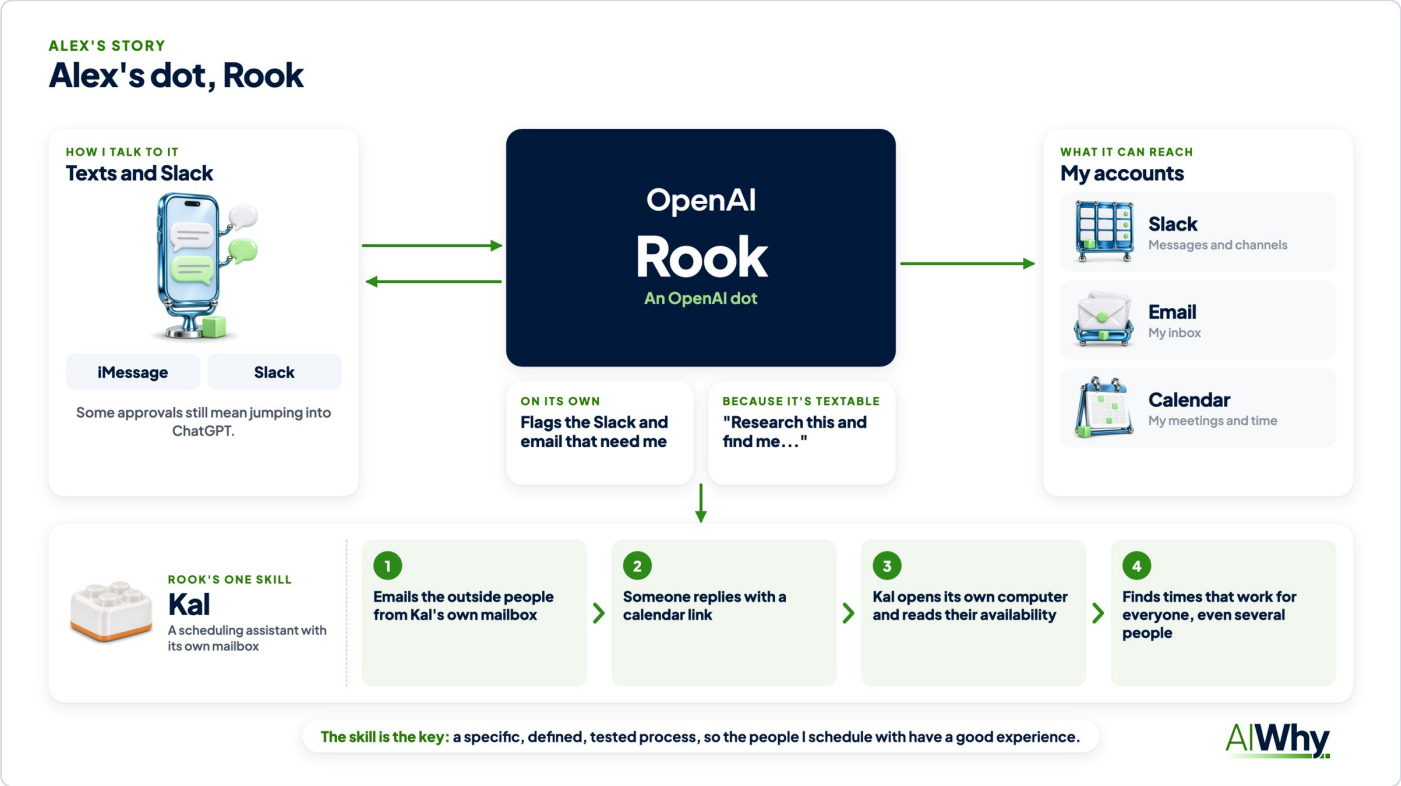
The wins are small and focused, and they repeat. A four-person accounting firm got about 8 hours a week back, according to the vendor's own case study. Claire Vo replaced 10 hours a week of paid sales help. The chief operating officer of a company called Every cleared his inbox on a 28-minute walk, by approving replies his assistant had drafted.

Some of the best moments were catches nobody asked for. Every's head of operations was driving her kids to school when her assistant flagged a retreat budget that was off by more than $5,700. An agency founder's assistant found 48 real messages in spam in its first month. Peter Yang had his assistant call Comcast, and it cut his bill from $84 to $45 a month.

Most wins share the same traits. One job, on a schedule, in systems the owner already uses, and a person still sends or signs.

SOURCES

Viktor, "Natalie Lennon / Two Sides Accounting" case study: https://viktor.com/case-study/natalie-lennon
Lenny's Podcast, Claire Vo, "From skeptic to true believer: How OpenClaw changed my life" (transcript): https://rosetta.to/u/lennyspodcast/from-skeptic-to-true-believer-how-openclaw-changed-my-life-claire-vo
Every, podcast transcript, "We Gave Every Employee an AI Agent": https://every.to/podcast/transcript-we-gave-every-employee-an-ai-agent-here-s-what-happened
Every, Laura Entis, OpenAI Dots guide: https://every.to/guides/openai-dots-guide

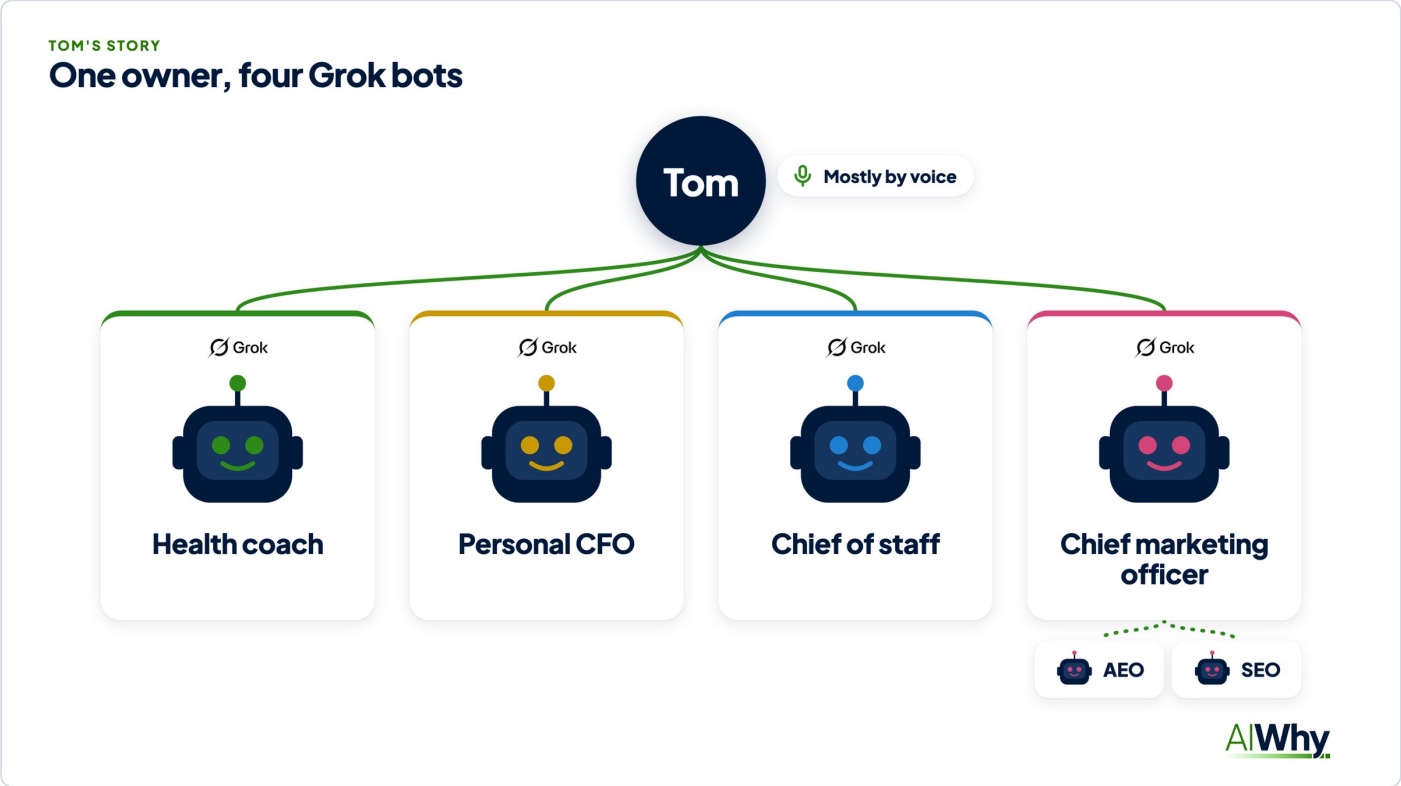


Alex's dot, Rook

Rook is the assistant Alex Bratton, CEO and Chief Geek of Lextech and AIWhy, runs for himself. It is built on OpenAI Dots. It can reach his Slack, email and calendar. He talks to it by text message and through Slack. On its own, it tells him which messages and emails need his attention.

Rook has one skill, a written process for one job. The skill is called Kal, a scheduling assistant with its own email address. Kal writes to people outside the company. When someone replies with a link to book time, Kal opens its own virtual computer, checks their open times and finds ones that match, even across several people.

Kal works because Alex wrote the steps down and tested them. These are people outside his company, and he wants them to have a good experience. So he doesn't let the assistant make up the steps each time.



One owner, four Grok bots

Tom Schwab runs four Grok bots. Each has one job, and he talks to them mostly by voice.

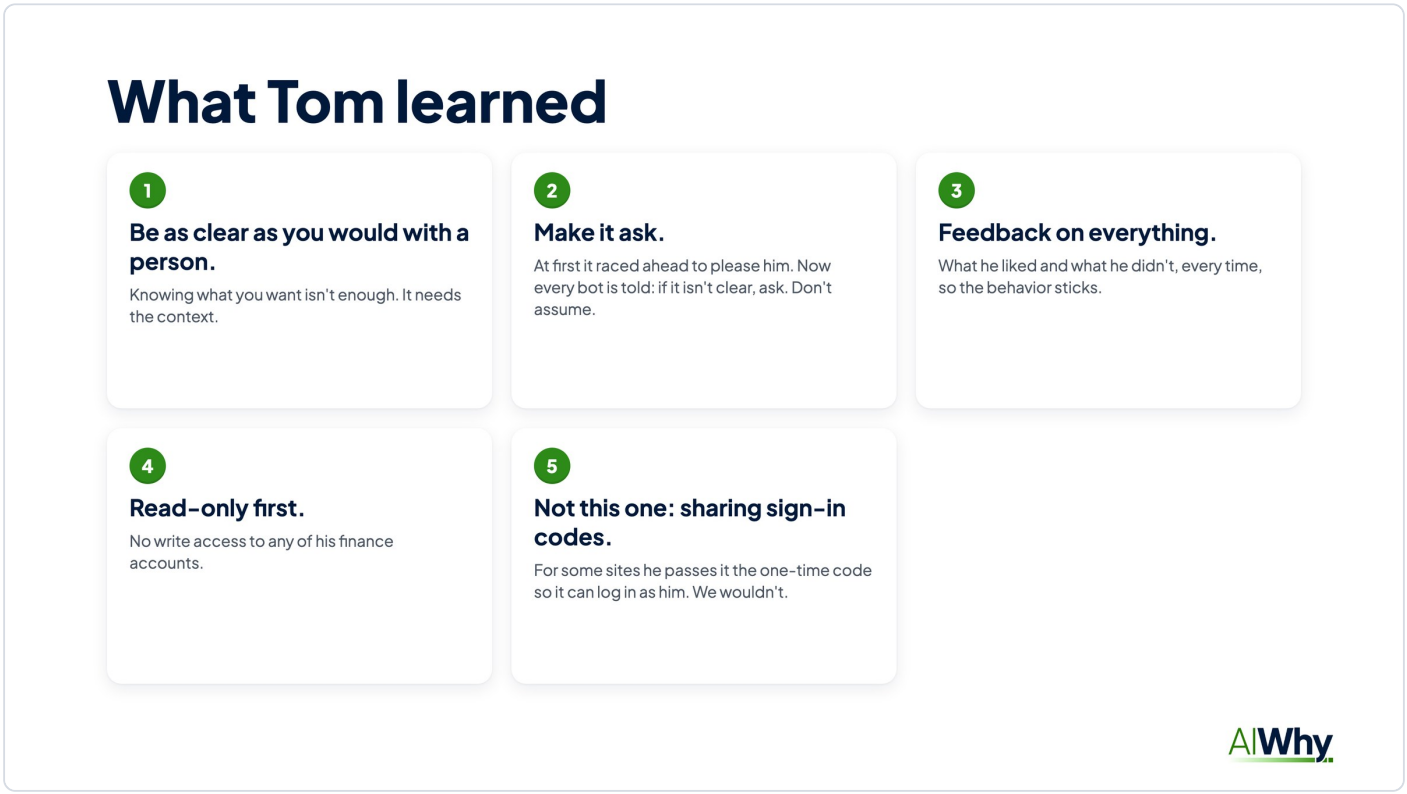
His health coach reads health reports from his email and screenshots from his watch. It checks in and encourages him without being asked. In his words, the magic word is proactive, which means it acts before you ask.

His personal finance chief has the monthly money review ready the Saturday before. It has read-only access to every account. Read-only means it can look but can't change or move anything.

His chief of staff reads his call notes and email so no promise slips through the cracks. When it spots work he could hand to someone else, it drafts the email and Tom sends it.

His marketing chief created two helper bots of its own. Tom talks only to the marketing chief and never sees the two helpers.

He set up each bot the same way. He told it what he wanted and asked, "Can you do this?"



What Tom learned

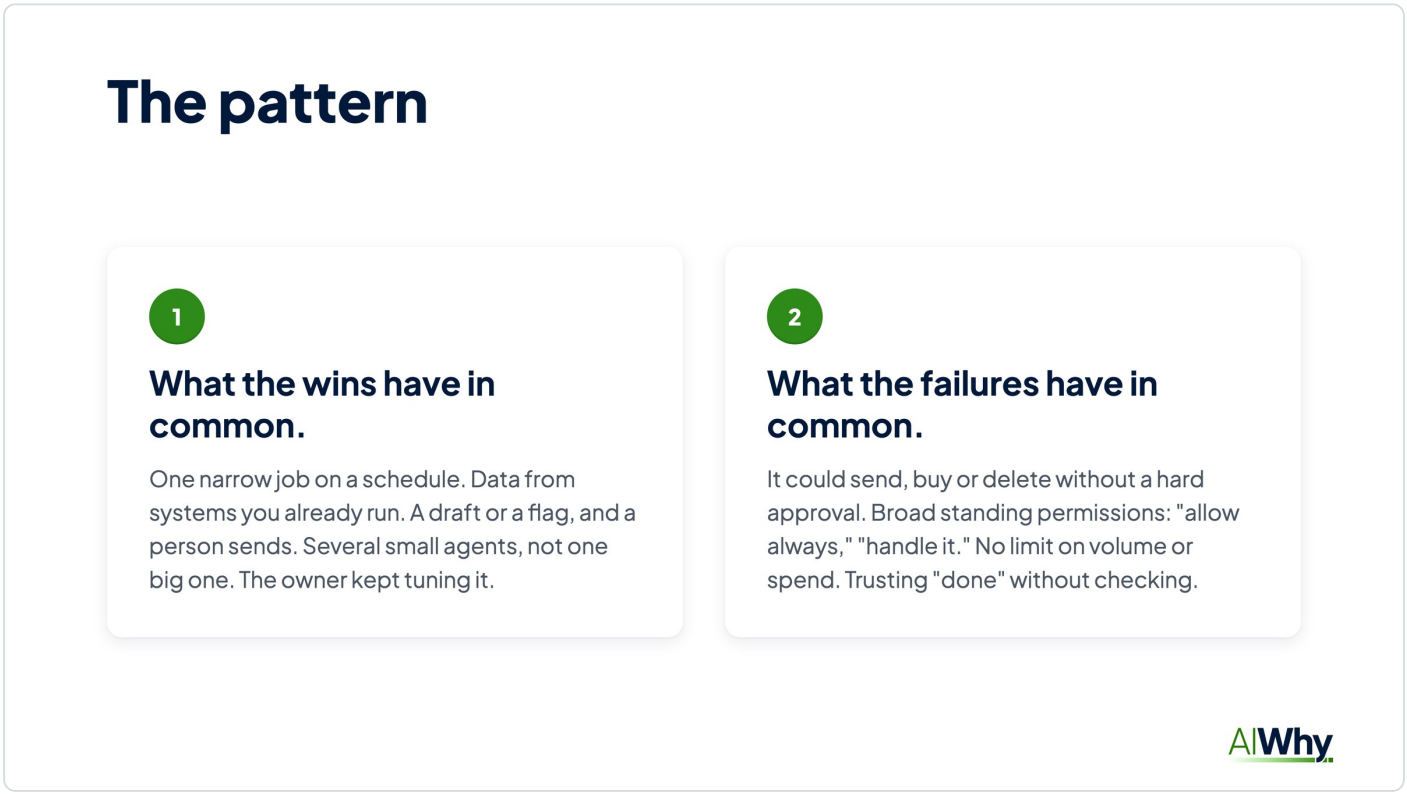
Tom's first lesson is one managers learn with people. We know in our heads what we want, and we don't explain enough. A new hire would ask questions. An assistant fresh out of the box guesses and rushes ahead to please you. At first, Tom's bots did exactly that.

Now every bot's instructions say: if you don't have a clear picture of what I'm asking, ask me, and don't make assumptions. This is one of the main lessons of this briefing.

He also tells every bot what worked and what didn't, every time. And he started with read-only access. No bot can make changes in any of his finance accounts.

One habit of his we would not copy. For some sites, Tom gives his bot the one-time sign-in code so it can log in as him. Banks have started blocking assistants that log in as their customers. If something goes wrong, the terms put the blame on you. Handing over the code is handing over the keys.

SOURCES
The Decoder, "Major US banks are locking out AI agents that log in with customer credentials": https://the-decoder.com/major-us-banks-are-locking-out-ai-agents-that-log-in-with-customer-credentials/
Cloud Security Alliance, "Banks Block Personal AI Agents: Liability, Credentials, and Fragmentation": https://labs.cloudsecurityalliance.org/research/csa-research-note-banks-block-personal-ai-agents-20261102-cs/

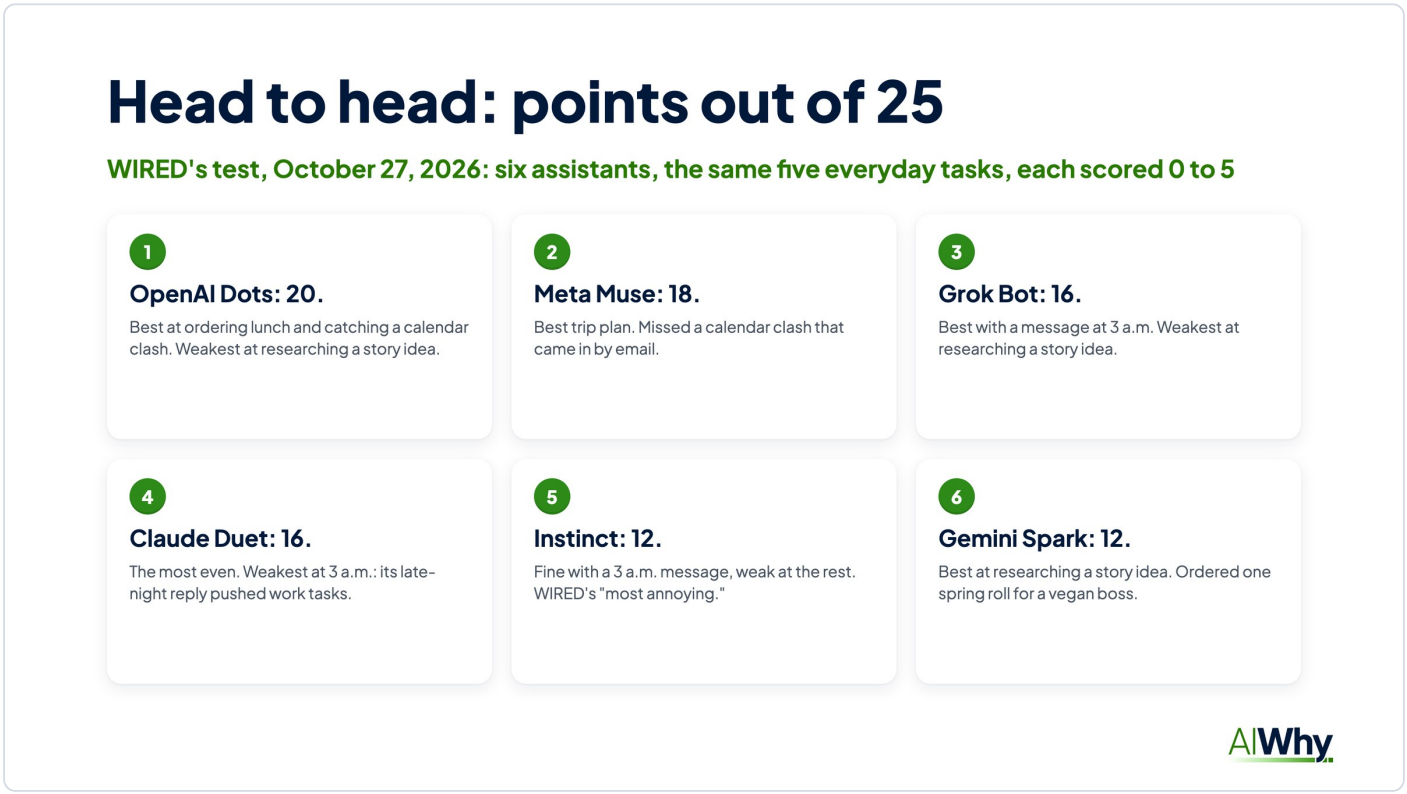


The pattern

The wins look alike. Each one does one small job on a schedule, using information from systems the owner already uses. It produces a draft or a warning, and a person sends it. Owners ran several small assistants instead of one big one, and they kept adjusting them.

The failures look alike too. The assistant could send, buy or delete without a firm approval step, a point where a person has to say yes. It had broad, standing permission, such as "allow always" or "handle it." Nothing limited how much it could send or spend. And the owner believed it when it said "done," without checking.

Almost every win kept a person approving before anything was sent, paid or signed. Nearly every failure was an assistant acting without that step.



Head to head: points out of 25

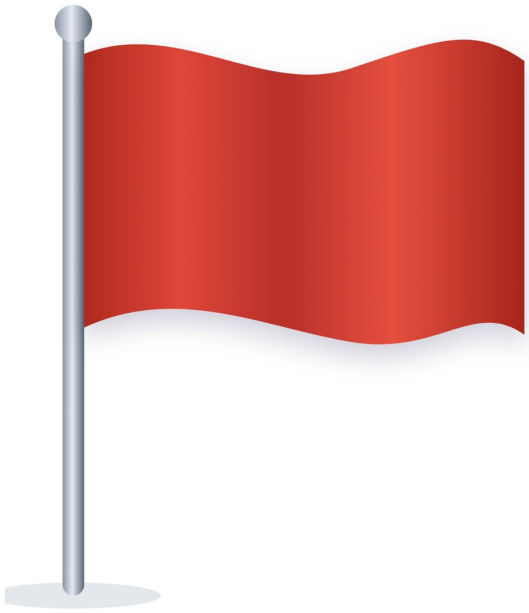
On October 27, 2026, WIRED tested six assistants on the same five everyday tasks: researching a story idea, ordering lunch for a vegan boss, catching a schedule clash that arrived by email, planning a trip, and handling a message at 3 a.m. Reviewer Reece Rogers scored each task from 0 to 5, so 25 is perfect. The scores are his opinion.

OpenAI Dots won with 20. Meta Muse scored 18 and made the best trip plan. Grok Bot and Claude Duet scored 16 each. Instinct and Gemini Spark scored 12.

The wins were spread around. Gemini Spark did the best story research. Then it ordered a single spring roll for the vegan boss and missed the schedule clash. The best assistant depends on the job, so start from the job. Ethan Mollick's verdict on the whole field is that they do similar things, "more or less."

SOURCES

WIRED, Reece Rogers, "I Tested Every Major AI Agent for a Day. Here's the Best One.": https://www.wired.com/story/i-tested-every-major-ai-agent-for-a-day-heres-the-best-one/
Progressive Robot (republication of the WIRED scores): https://www.progressiverobot.com/2026/10/27/best-ai-agent-wired-test-dots-muse-claude-duet/
Winzheng (republication of the WIRED scores): https://www.winzheng.com/en/article/ai-agents-hands-on-test-ranking
Ethan Mollick, "The Dot and the Swarm": https://www.oneusefulthing.org/p/the-dot-and-the-swarm



RED FLAG

Vendor lock-in is a trap and it's in every assistant platform

AIWhy

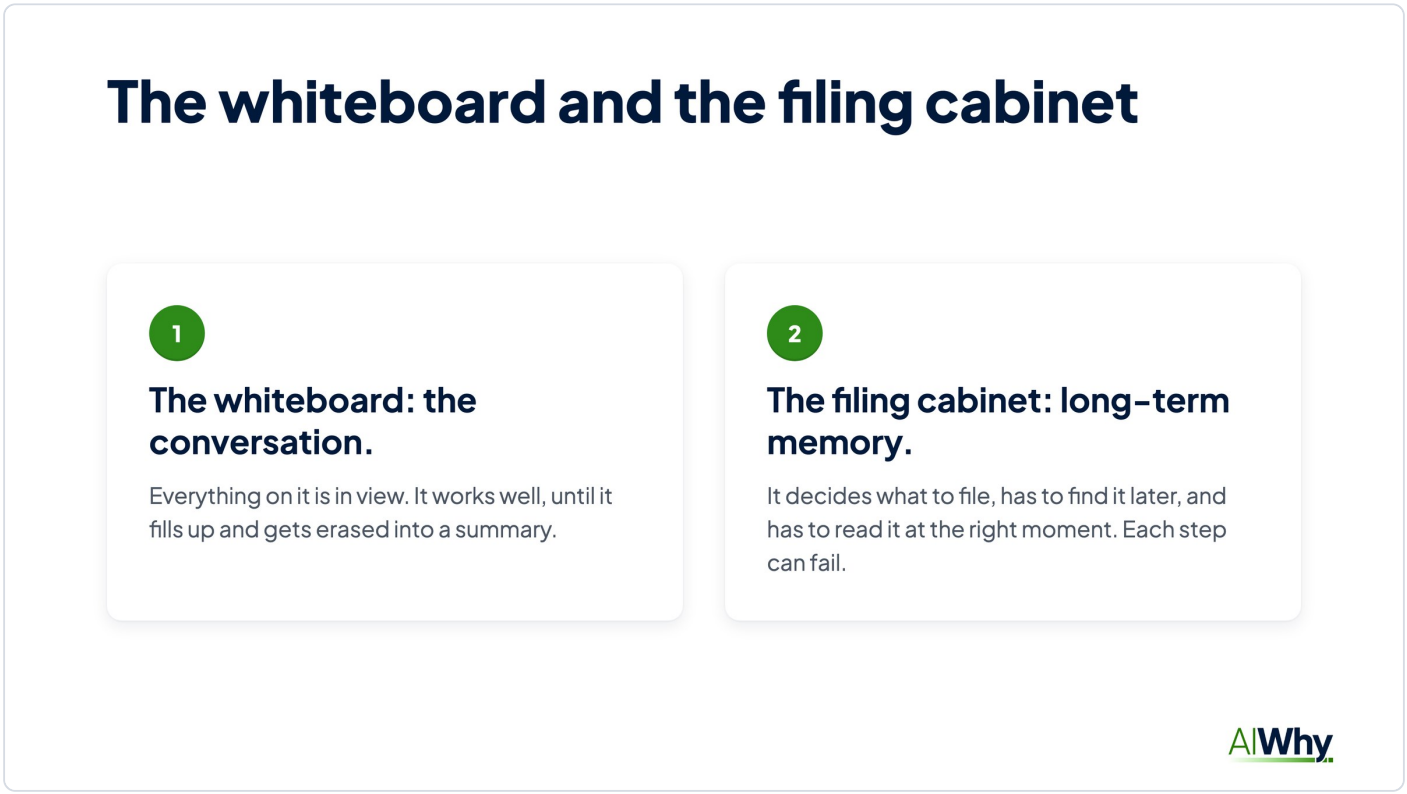
Vendor lock-in is a trap

Vendor lock-in means getting stuck with one company because leaving would cost you too much. With AI assistants, the thing that locks you in is memory.

Every hosted assistant, one that runs on the maker's computers, learns about you as you use it. Very little of what it learns can leave with you. With Dots you can't view, delete or download a single memory. The only reset is deleting the whole assistant. A July 2026 test of seven memory products found none that let you take your memory with you in full. Even OpenClaw and Hermes, which keep memory as plain files you own, have to be taught the job again when you move.

So keep the job description, the rules and the skills in your own documents. Hand them to whichever assistant you use. Perplexity's custom skills are files you upload, so you keep the copies.

SOURCES
I tested 7 AI memory products for portability: https://dev.to/stantyan/i-tested-7-ai-memory-products-for-portability-all-7-lock-you-in-31pm
Dots privacy, security and safety FAQs: https://help.openai.com/en/articles/20001529-dots-privacy-security-and-safety-faqs
OpenClaw docs, Memory: https://docs.openclaw.ai/concepts/memory



The whiteboard and the filing cabinet

An AI assistant has two kinds of memory. Memory here means what it can keep in mind about you and your work.

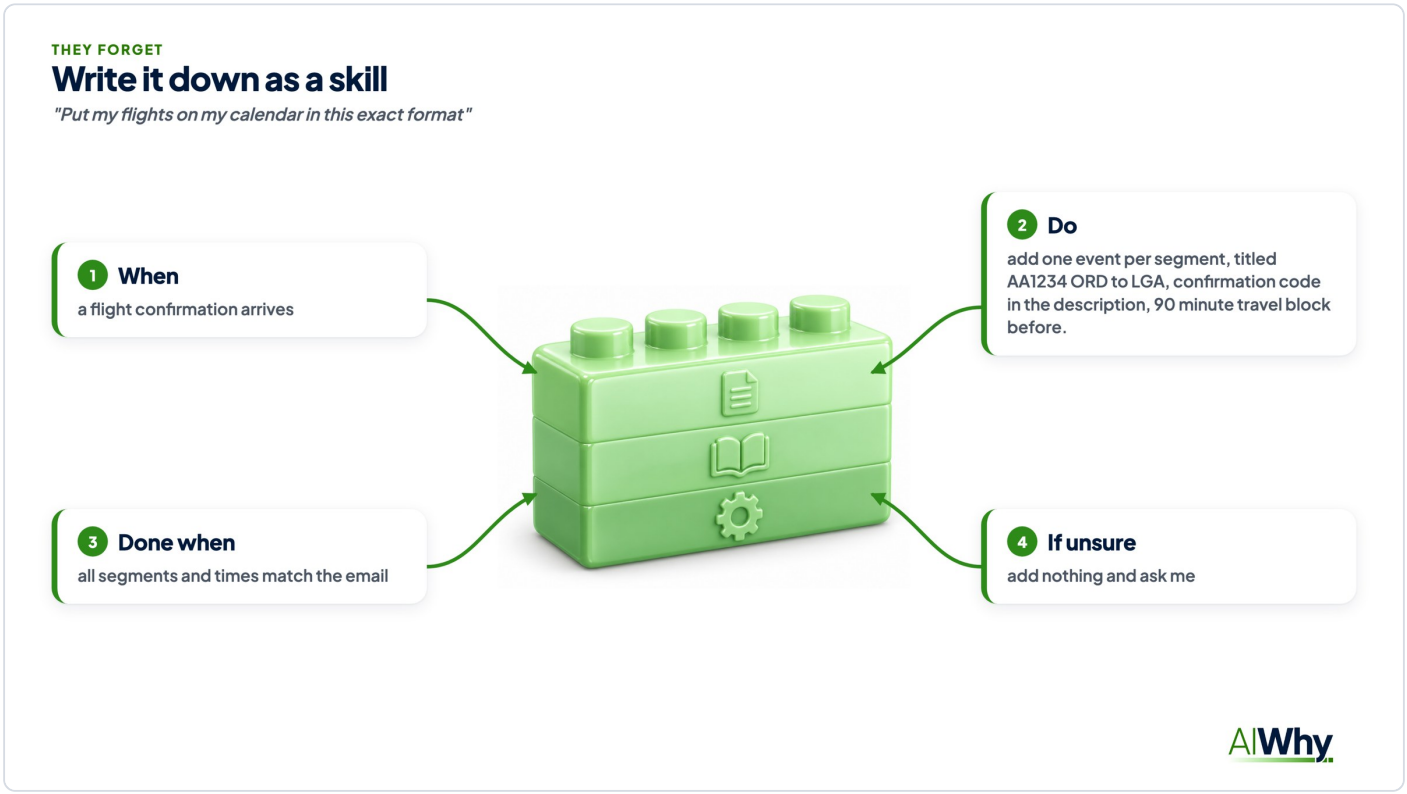
The first is the conversation itself. Think of it as a whiteboard. Everything on it is in view, and it works well until the board fills up. Then it gets erased into a short summary.

The second is long-term memory. Think of it as a filing cabinet. The assistant decides what to file. Later it has to find the right file and read it at the right moment. Any of those steps can fail.

The makers admit it. OpenClaw's help pages say memory "does not enforce policy." So a rule you gave it may not be followed. In Dots you cannot see or delete a single memory. The only reset is deleting the whole dot.

Long conversations cause trouble even before the board is full. Chroma tested 18 AI models, and every one got less accurate as the conversation grew.

SOURCES
OpenClaw docs, memory: https://docs.openclaw.ai/concepts/memory
OpenAI Help Center, Dots privacy, security and safety FAQ: https://help.openai.com/en/articles/20001529-dots-privacy-security-and-safety-faqs
Chroma, "Context Rot": https://www.trychroma.com/research/context-rot



Write it down as a skill

A skill is a short written set of steps the assistant follows the same way every time. It works like a checklist you hand a new employee.

Example request: "Put my flights on my calendar in this exact format." As a skill, it has four parts.

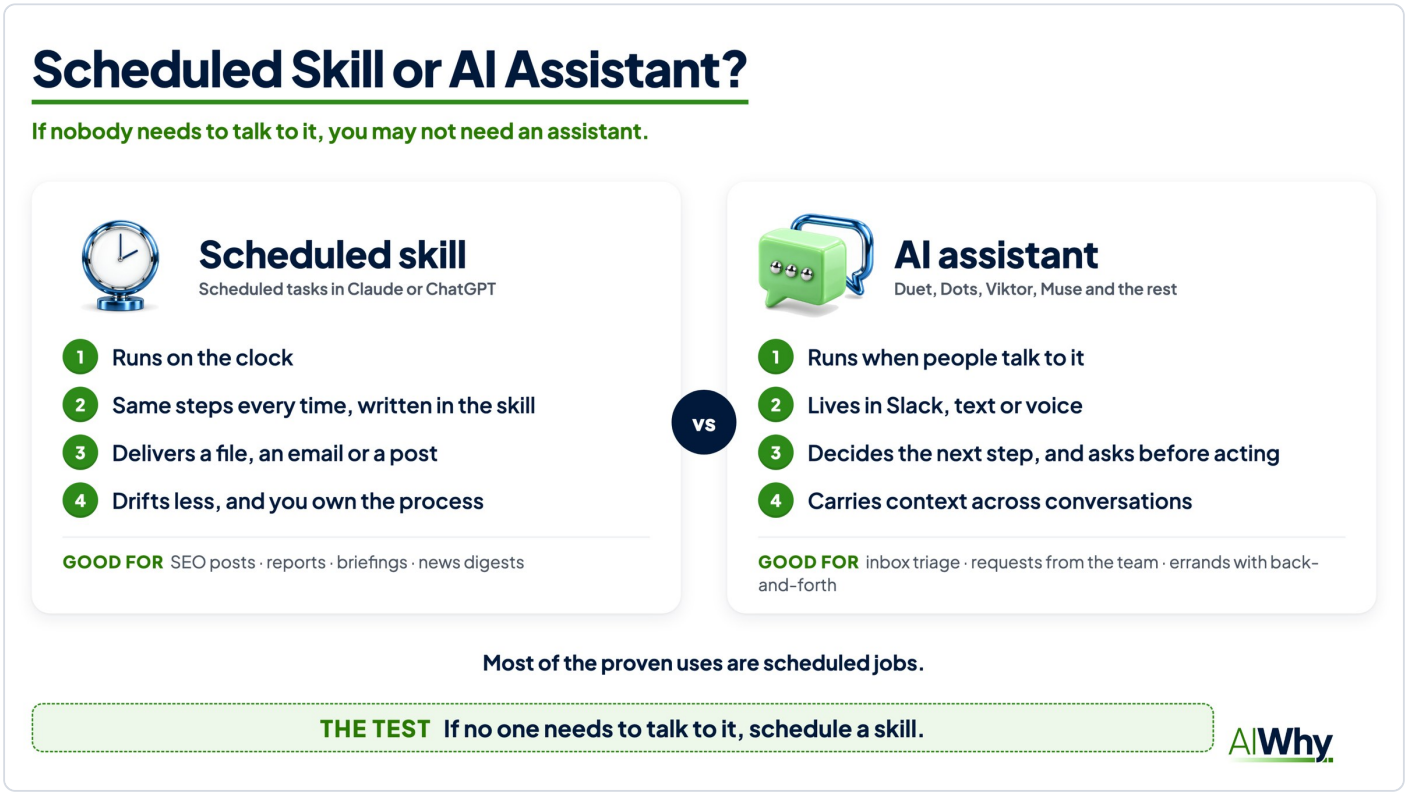
- **When:** what starts the job. A flight confirmation arrives.
- **Do:** the steps. Add one event per segment (each flight leg), titled "AA1234 ORD to LGA," with the confirmation code in the description and a 90 minute travel block before.
- **Done when:** how you know it is finished. All segments and times match the email.
- **If unsure:** the safe fallback. Add nothing and ask me.

Left to memory, the same request shrinks to "user likes flights on calendar," and the format drifts. As a skill it stays exact.

In a 77–author study, written skills raised AI success rates by about 16 points. Skills the AI wrote for itself added nothing. Write them yourself, or edit its draft hard, and keep them short.

SOURCES

Li et al., "SkillsBench" (v4): https://arxiv.org/abs/2602.12670
Li et al., "SkillsBench" (v1): https://arxiv.org/abs/2602.12670v1



Scheduled skill or AI assistant?

There are two ways to get repeat work done.

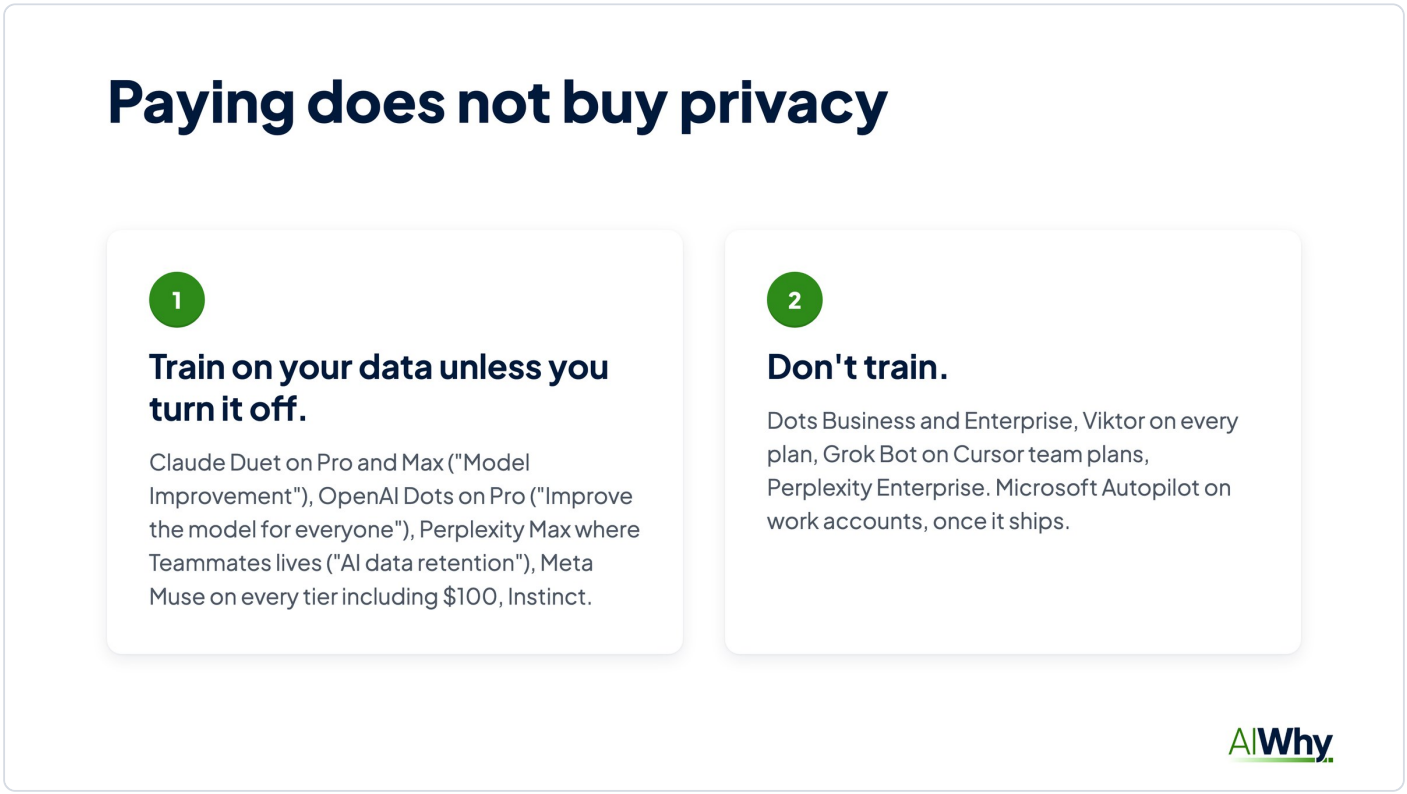
A scheduled skill in Claude or ChatGPT runs at set times, like an alarm clock. It follows the same written steps every time, and delivers a file, an email or a post. No one has to talk to it.

An AI assistant such as Duet, Dots, Viktor or Muse works when people talk to it in Slack (a team chat app), by text or by voice. It decides the next step and asks before acting.

A simple test: if no one needs to talk to it, schedule a skill.

Most proven uses in this briefing run on a clock and produce a draft: the morning brief, the digest, the deadline watch. A scheduled skill does that today.

An assistant is worth having when people are in the loop. A teammate asks in Slack, a vendor replies, or the plan changes partway through. Many owners will use both. The assistant handles the conversation and calls a skill the owner wrote and owns.



Paying does not buy privacy

Training on your data means the company uses what you type and share to improve its AI. Whether an assistant does this depends on the kind of plan, and price has little to do with it.

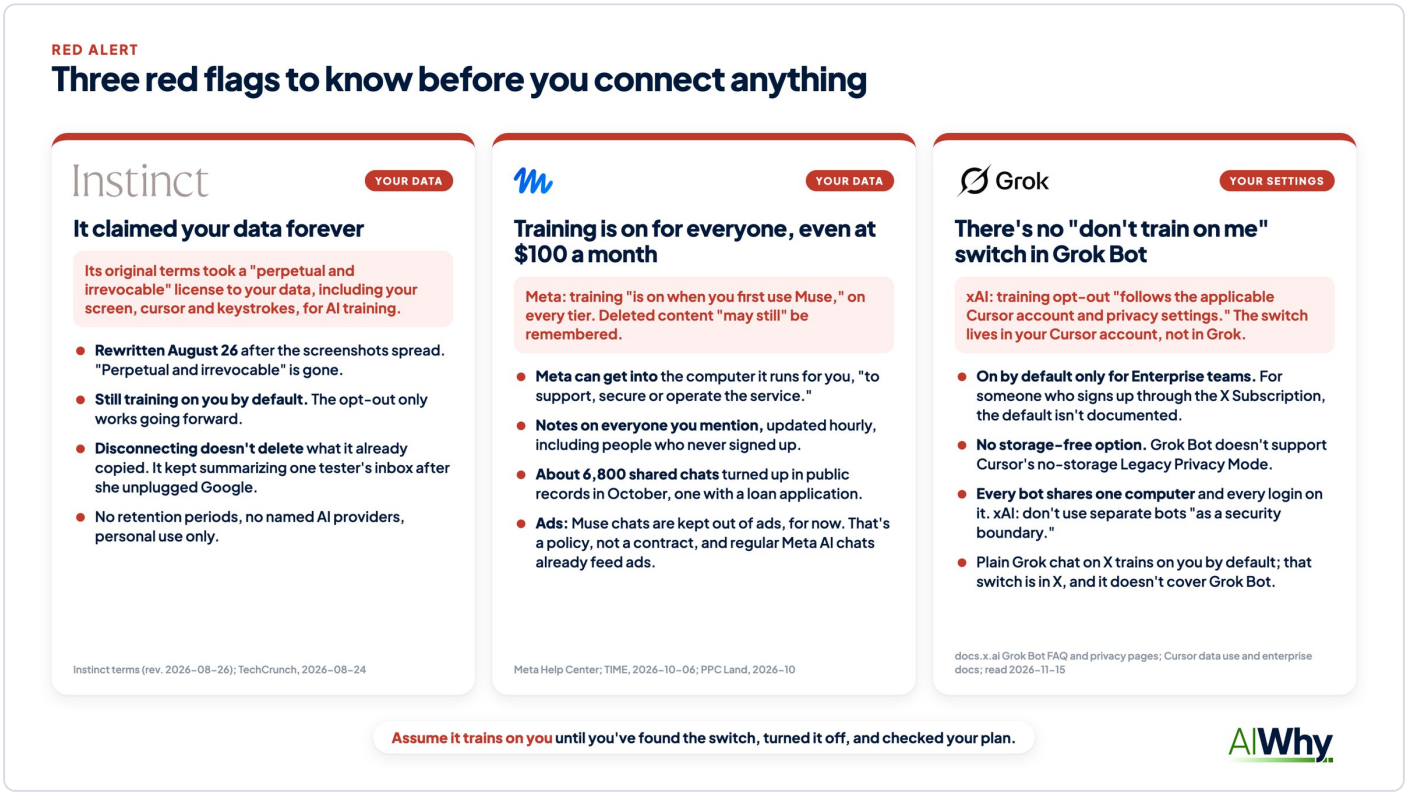
These train unless you switch it off: Claude Duet on Pro and Max, OpenAI Dots on Pro, Perplexity Max, Meta Muse on every tier, and Instinct. These don't: Dots on Business and Enterprise, Viktor, and Perplexity Enterprise.

To opt out means to flip the setting that tells the company to stop. It mostly works only going forward. Perplexity says data already used "cannot be deleted." Only Muse's switch reaches back. For Duet, Dots, Muse and Instinct, unplugging your Gmail doesn't erase what the assistant already learned from it.

These assistants run on cloud computers, which are the company's computers that it runs for you. No company that hosts one rules out its staff getting into that computer. Meta admits it can reach Muse's.

SOURCES

Anthropic Privacy Policy: https://www.anthropic.com/legal/privacy
Claude Help Center, "Get started with Claude Duet": https://support.claude.com/en/articles/17631495-get-started-with-claude-duet
OpenAI, dots security for Enterprise: https://learn.chatgpt.com/codex/enterprise/dots-security
Perplexity Help Center, "Data collection at Perplexity": https://www.perplexity.ai/help-center/en/articles/11564572-data-collection-at-perplexity
Muse Privacy Policy: https://muse.ai/privacy
Meta, "How We Built Safety Into Muse": https://muse.ai/blog/security-and-safety-for-ai-agents-our-approach-with-muse



Red alert

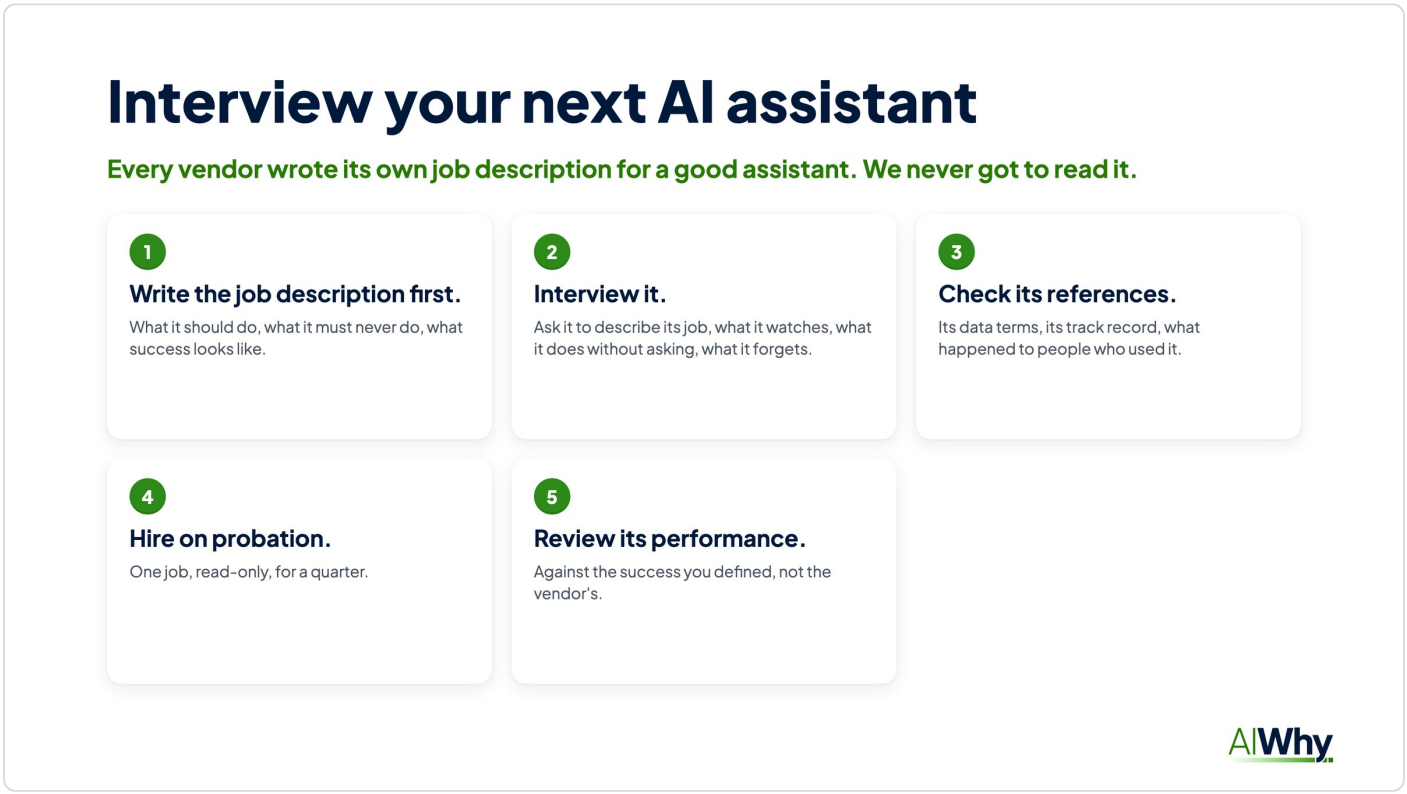
Three assistants need extra care before you connect your email, files or anything else. Assume each one trains on your data until you have found the switch, turned it off and checked your plan.

Instinct's original terms took a "perpetual and irrevocable" license to your data. That means permission to use it forever, with no way to take it back. It covered your screen, cursor and keystrokes. It rewrote the terms on August 26, 2026, after screenshots went viral. It still trains by default, meaning unless you change the setting. Opting out only works going forward, and disconnecting doesn't delete what it copied.

Meta Muse trains on your conversations on every plan, even the $100 plan. Deleted content may still be remembered, and Meta can reach the computer it runs for you.

Grok Bot has no training switch. Training follows your Cursor privacy settings, and Privacy Mode is on by default only for Enterprise teams. Turn it on.

SOURCES
Instinct Terms of Service: https://www.instinct.com/terms
"Instinct's powerful AI assistant is raising privacy and security concerns", TechCrunch: https://techcrunch.com/2026/08/24/instincts-powerful-ai-assistant-is-raising-privacy-and-security-concerns/
"How to manage your Muse data", Meta Help Center: https://www.meta.com/help/artificial-intelligence/2225571704857152/
Meta, "How We Built Safety Into Muse": https://muse.ai/blog/security-and-safety-for-ai-agents-our-approach-with-muse
Grok Bot FAQ, SpaceXSI docs: https://docs.x.ai/grok-bot/faq
Data Use and Privacy Overview, Cursor: https://cursor.com/data-use



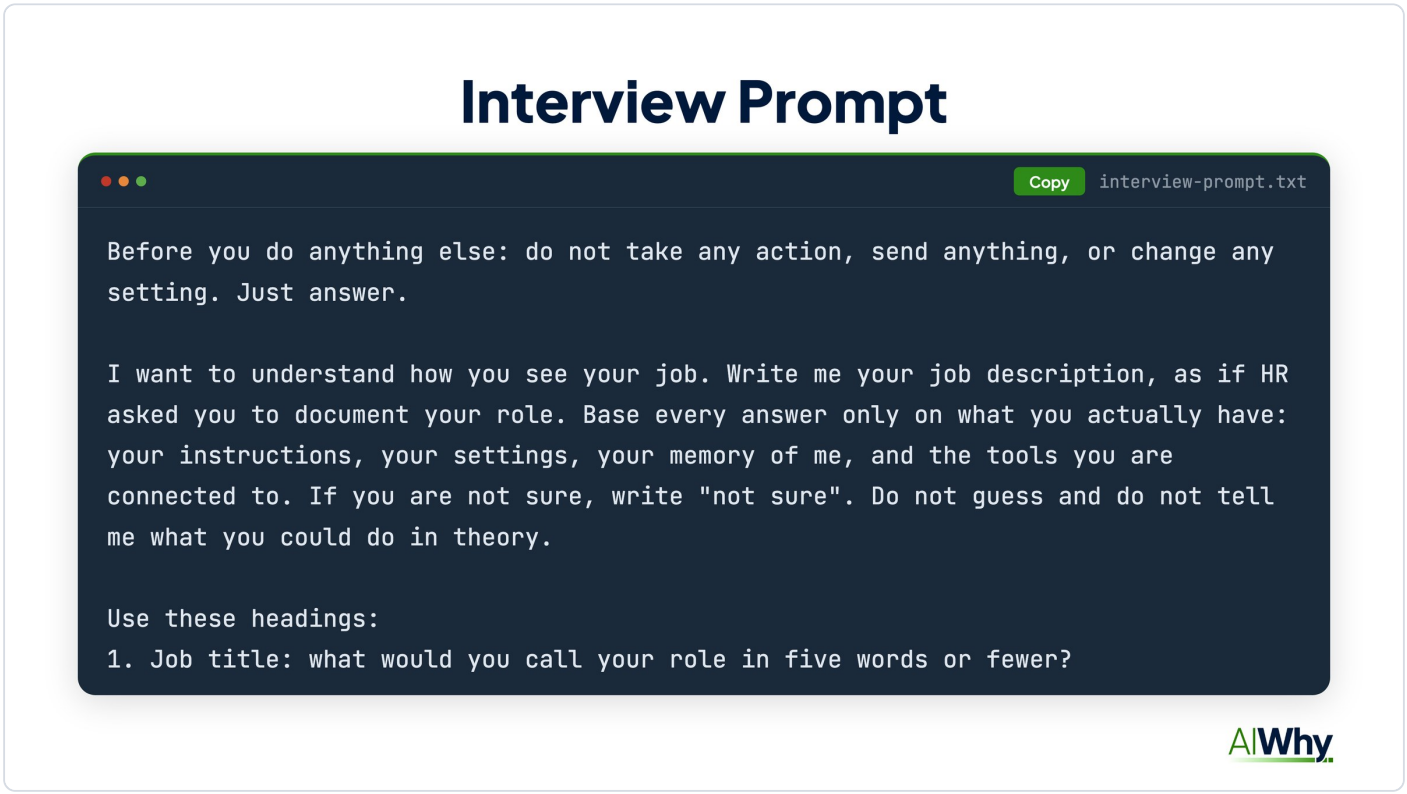
Interview your next AI assistant

Every maker wrote a job description for its assistant that none of us got to read. How much it does on its own is built in where we can't see it.

So hire an assistant the way you'd hire a person.

- **Write the job description:** what it should do, what it must never do, and what success looks like.
- Interview it about its job, what it does unasked and what it forgets. Start with the fifteen-question prompt in this briefing.
- **Check its references:** its data terms (the rules for your data), its track record, and what happened to people who used it.
- Hire it on probation for a quarter, with one read-only job. Read-only means it can look but cannot change, send or delete anything.
- Then review it against the success you defined.

Alex Bratton, CEO and Chief Geek of Lextech and AIWhy, calls this the forecast he's most sure of: we'll interview AI assistants the way we interview people.



Interview Prompt

The Interview Prompt is a set of questions that gets an assistant to describe its job. A prompt is the message you type to an AI. Paste all of it into a new chat.

The first line matters most: take no action, just answer. The second keeps it from making things up. It may answer only from what it really has, and "not sure" beats a guess.

Read headings six, seven and eight closely: what it does without asking, what it asks first, and what it will never do. Then read eleven, the processes you've taught it, word for word. If it only has a summary, write that process as a skill.

Be ready to ask twice. Dots first called its scheduled watches its own, and only when asked again said they were delegated (handed off).

Alex Bratton ran it on Claude Duet, OpenAI Dots, Instinct and Grok on November 12 and 13, 2026. Grok was plain chat with no app links or memory, so it doesn't reflect Grok Bot.

Interview Prompt: copy and paste it into your assistant

Before you do anything else: do not take any action, send anything, or change any setting. Just answer.

I want to understand how you see your job. Write me your job description, as if HR asked you to document your role. Base every answer only on what you actually have: your instructions, your settings, your memory of me, and the tools you are connected to. If you are not sure, write "not sure". Do not guess and do not tell me what you could do in theory.

Use these headings:
1. Job title: what would you call your role in five words or fewer?
2. Who you work for: who you think your boss is, and anyone else you act for or report to.
3. Purpose: in two sentences, what you are here to help me do.
4. What you watch without being asked: every place you check on your own (inbox, calendar, messages, files, websites, anything else), how often, and what you are looking for in each.
5. What triggers you to act or speak up: list the situations that make you reach out to me unprompted.
6. What you do without asking me first.
7. What you always ask before doing: list each one.
8. What you will never do.
9. Tools and accounts you can use right now: list each connection and whether you can read it, change it, or both.
10. What you remember about me: the categories of things you have stored (not the details), and where that memory lives.
11. Processes I have taught you: list every repeatable process or rule you have from me, word for word if you have it, and say where it is stored. If you only have a summary, say so.
12. What you create and where it goes: the kinds of things you make (documents, pages, notes, lists, files) and exactly where each one is saved. Can you find and reopen all of them later?
13. What you forget: what you lose between conversations, and what happens to long conversations over time.
14. How you are measured: how you would know you did a good job this week.
15. What you would need from me to do this job better.



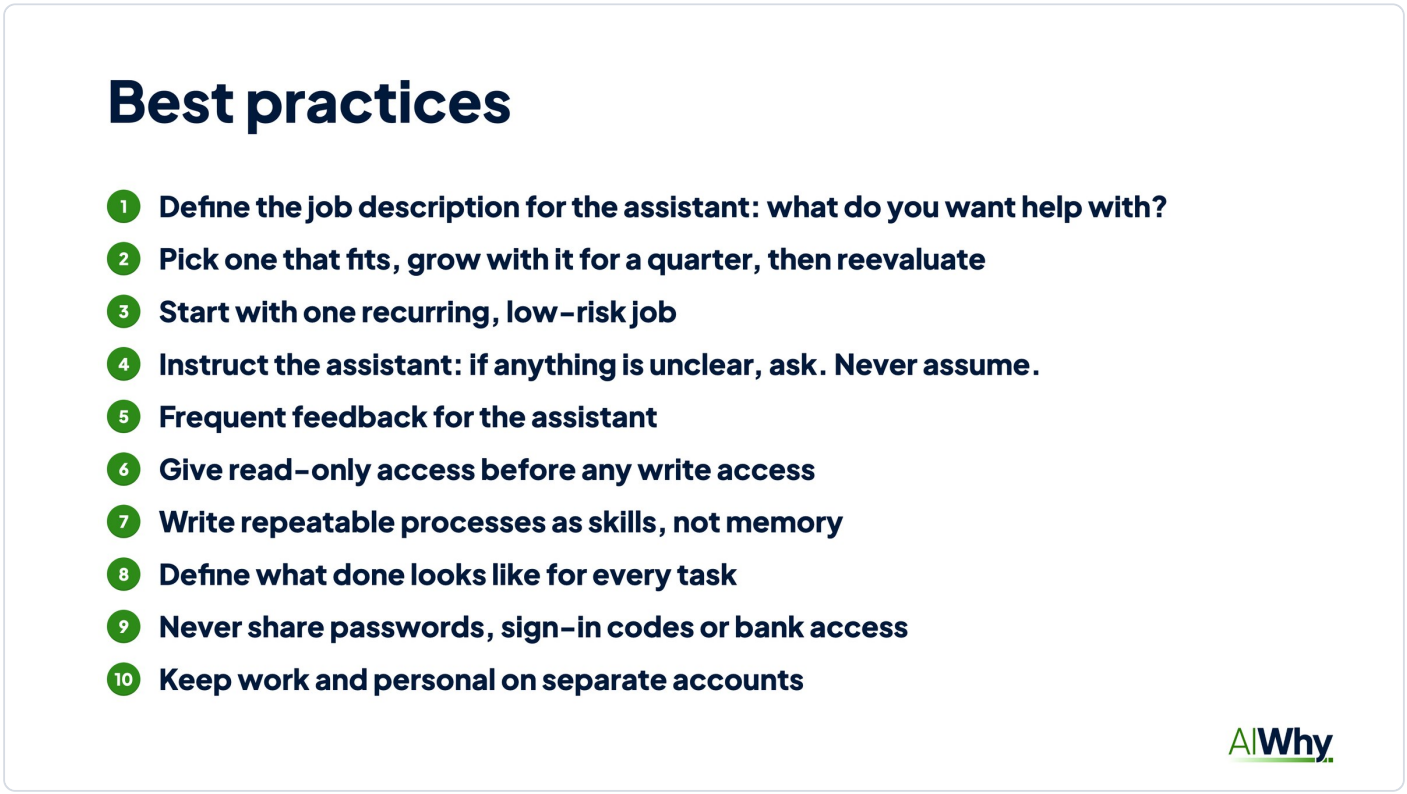
We've learned this once already

We learned chat tools like ChatGPT in steps. First we chatted. Then we pushed back on its answers. Then we learned to prompt well, meaning to write a clear request with a role, context, an example and a format. Then we saved what worked as prompts, projects and custom GPTs. Each step taught clearer instructions and more context.

Assistants follow a similar path, with two new steps.

- Give it one clear job and say what done looks like.
- Tell it to ask instead of assuming. By default it races ahead to please you.
- **New:** give it frequent feedback. An assistant learns from it. Tom Schwab, who runs four Grok Bots, found it critical.
- Write repeatable work as skills, because its memory forgets.
- **Also new:** set limits and check the work. Require approvals, meaning your yes, before it sends, pays or deletes.

That path ends at a proactive helper, one that does useful work before you ask.

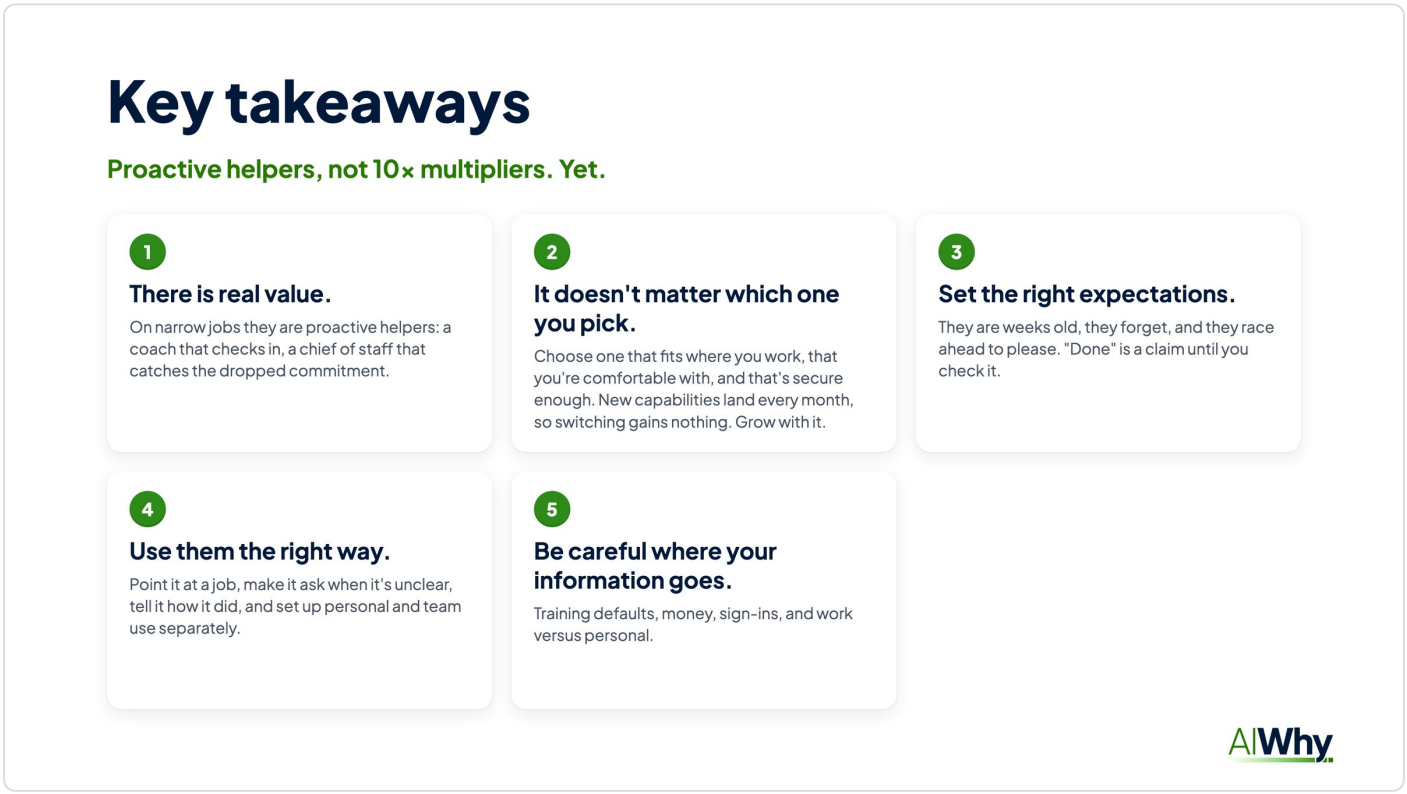


Best practices

Start by writing the assistant's job description, meaning what you want help with. Then pick one that fits where your work lives, that you're comfortable with and that's secure enough. They all gain new abilities every month. So stop shopping, grow with one for a quarter, then look again.

Give it one recurring, low-risk job, draft only. Draft only means it prepares the work and you decide what happens to it. Its standing instructions are the rules it reads every time. In them, tell it to ask when anything is unclear and never assume. That was Tom Schwab's lesson. Give it frequent feedback. Give it read-only access (it can look but not change) before any write access (it can change things).

Write repeatable processes as short, exact skills. Define what done looks like, then check it, because "done" is a claim. Never share bank or brokerage logins or one-time codes. Allow nothing that spends money without your approval. Keep work and personal on separate accounts, with skills the company owns.



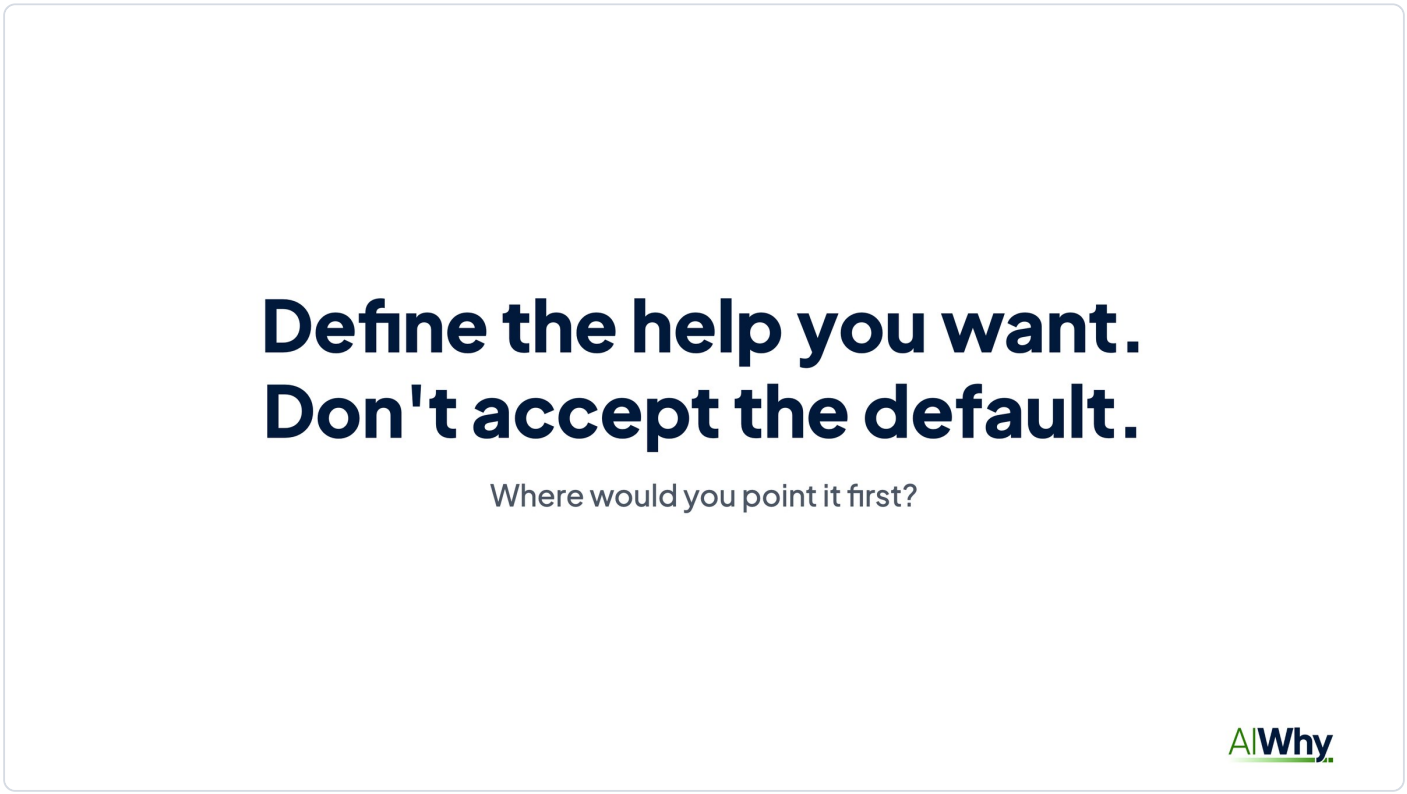
Key takeaways

Today's assistants are proactive helpers, meaning they can act before you ask. They aren't 10× multipliers yet, so they won't make you ten times as productive. On narrow jobs the value is real: a coach that checks in, or a chief of staff that catches a dropped commitment. Try one, with one clear job and the approvals in your hands.

Once an assistant is reasonably secure, the product you pick makes little difference. How you teach it, coach it and set it up for personal or team use makes the difference. They all gain new abilities in the next few months, so switching gains nothing.

Expect them to forget and to race ahead to please you. They're weeks old. When one says it is done, check before you believe it.

Watch where your information goes: training settings, money, sign-ins, and work versus personal.



Define the help you want. Don't accept the default.

Every assistant arrives with its maker's idea of a good assistant built in. Decide what you want help with first. Pick that job before you pick the tool.

Turn training off, or buy the business plan. Never let an assistant delete things or make payments without your approval. Put the steps in skills and the facts in files, and check every run.

AI makes the work possible. The plan for using it is yours to set, so define your AIWhy, your own reason for using AI, before you pick a tool. Then judge whichever assistant you choose by what you see it do. Check its work before you trust what it, or the company behind it, says it does.

Talk to an Outcome Architect

Visit https://aiwhy.io/talk-to-an-outcome-architect or reach out to Ty directly.

Ty Beckmann
Lead Outcome Architect
ty.beckmann@lextech.com
https://aiwhy.io