AIWhy Briefing

AI Assistants

Define your AIWhy before you pick a tool

Briefing as of November 17, 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™ 5
"Which AI assistant should I use?" is the wrong question 7

01 Ground rules 9
Anatomy of an AI assistant 10

02 The field is exploding 12
From a category of one to crowded 13
Every month, something new 15

03 Who's who 17
How much assistant comes built in? 18
What they say they are for 20
What they all say about their job 22
Same question, four very different jobs 24

04 What to expect 26
What's actually working 27
Alex's dot, Rook 29
One owner, four Grok bots 31
What Tom learned 33
What breaks 35
The pattern 37
Head to head: points out of 25 39
Vendor lock-in is a trap 41

05 They forget 43
The whiteboard and the filing cabinet 44
Write it down as a skill 46
Scheduled skill or AI assistant? 48

06 Your data, your money, your company 50
Paying does not buy privacy 51
Red alert 53
Money is where it gets real 55
Keep work and personal apart 57

07 Using them well 59
Interview your next AI assistant 60
Interview Prompt 62
We've learned this once already 64
Best practices 66
Key takeaways 68
Define the help you want. Don't accept the default. 70

08 Appendix: platform by platform 72
Claude Duet 73
OpenAI Dots 75
Meta Muse 77
Instinct 79
Grok Bot 81
n8n Nod 83
Viktor 85
Perplexity Teammates 87
Microsoft Copilot Autopilot 89
OpenClaw 91
Hermes Agent 93

AIWhy

AI Assistants: Nov 2026 reality check

Define your AIWhy before you pick a tool

Claude · OpenAI · Muse · Grok · perplexity · Copilot

n8n · Instinct · viktor · OpenClaw · Hermes



AIWhy

AI Assistants: Nov 2026 reality check

This is AIWhy's look at always-on AI assistants as of mid-November 2026. In the ten weeks before it, Anthropic, OpenAI, Meta, xAI, Perplexity and n8n all shipped an assistant that works in the background on your behalf. Microsoft and Google are close behind.

Everyone wants to know which one to use. Our advice is to define your AIWhy first, then pick a tool. We cover Claude, OpenAI, Meta Muse, Grok, Perplexity, Microsoft Copilot, n8n, Instinct, Viktor, OpenClaw and Hermes, along with what works, what breaks, two owners' own setups, and how to interview an assistant before you hire it.

Most of these products are weeks old and their terms change monthly. Check any number here against the vendor's current pages 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™

Stop renting your competitive advantage

1. **Your know-how should be an asset.**
   Today it lives in heads and rented software.
2. **Make it run.**
   Runnable IP™: your knowledge, written so it runs.
3. **AI is the software.**
   Live it. Tune it. It fits how you really work.
4. **Build enterprise value.**
   Owned. Verifiable. Transferable. Compounding.
5. **Grow with the team you have.**
   More clients, bigger engagements, same team.



Alex Bratton

AIWhy

Capability Equity™

AIWhy works from a method called Capability Equity. Its core idea: stop renting your competitive advantage. It rests on five ideas.

Your know-how should be an asset, and today it lives in people's heads and in rented software. Make it run: Runnable IP™ is your knowledge, written so it runs. AI is the software: you live it and tune it, so it fits how you really work instead of how the software made you work. Build enterprise value that is owned, verifiable, transferable and compounding. Grow with the team you have, taking on more clients and bigger engagements with the same people.

For AI assistants, that 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

The answer expires almost immediately. Instead ask why you want one.

1

The tool question.

Which one is best this week? They are weeks old, they change every month, and under the hood they are mostly the same machine.

2

The AIWhy question.

Where would I point it, and what must never happen without me? The answer outlasts every product name.

AIWhy

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

The tool question asks which assistant is best this week, and any answer expires almost at once. These products are weeks old, they change every month, and under the hood they are mostly the same machine.

The AIWhy question asks why you want one at all. Where would you point it, and what must never happen without you? That answer outlasts every product name.

The wins in this briefing all start from a job: a morning brief, a lead sweep, a weekly timesheet check. Nobody in the public record succeeded with "run my life."

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

01

Ground rules

AIWhy

GROUND RULES

Anatomy of an AI assistant

CHANNELS

How you reach it



App | Text, iMessage

Slack, Teams | Voice, phone

YOUR COMPUTER

Your own files



Through a desktop app, on some assistants

Opt-in

RUNS IN THE VENDOR'S CLOUD



The brain

AI model in the vendor's data center



Its own computer

A tool it uses: browser, files



Approval gate

Asks before it sends, pays or deletes



Memory

Notes it keeps about you



Skills

Your written playbooks



Schedule

Wakes it to check in

CONNECTORS

Keys to your accounts



Email

Read, draft, send



Calendar

Meetings and time



CRM

Contacts and deals



Work tools

Slack, drives, tasks, bank

The brain thinks in the vendor's data center. Its computer is a tool it uses.

AIWhy

Anatomy of an AI assistant

An assistant has nine parts, and we use the same words for them throughout.

- **Channels:** how you reach it. The app, text, Slack or Teams, voice.
- **Brain:** the AI model, rented from a model company and running in its data center.
- **Its own computer:** a virtual computer with a browser and files, which the brain uses to click through websites as you. It is a tool. The thinking happens in the data center.
- **Memory:** the notes it keeps about you.
- **Skills:** the written playbooks you give it.
- **Schedule:** the alarm clock that wakes it to check your inbox and calendar. This is what "proactive" means.
- **Approval gate:** the check before it sends, pays or deletes. On most products a second AI decides when to ask.
- **Connectors:** the keys to your email, calendar, CRM and the rest.
- **Your computer:** four of the assistants (Dots, Grok Bot, Muse and Perplexity) can reach files on your own machine through a desktop 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/

02

The field is exploding

AIWhy

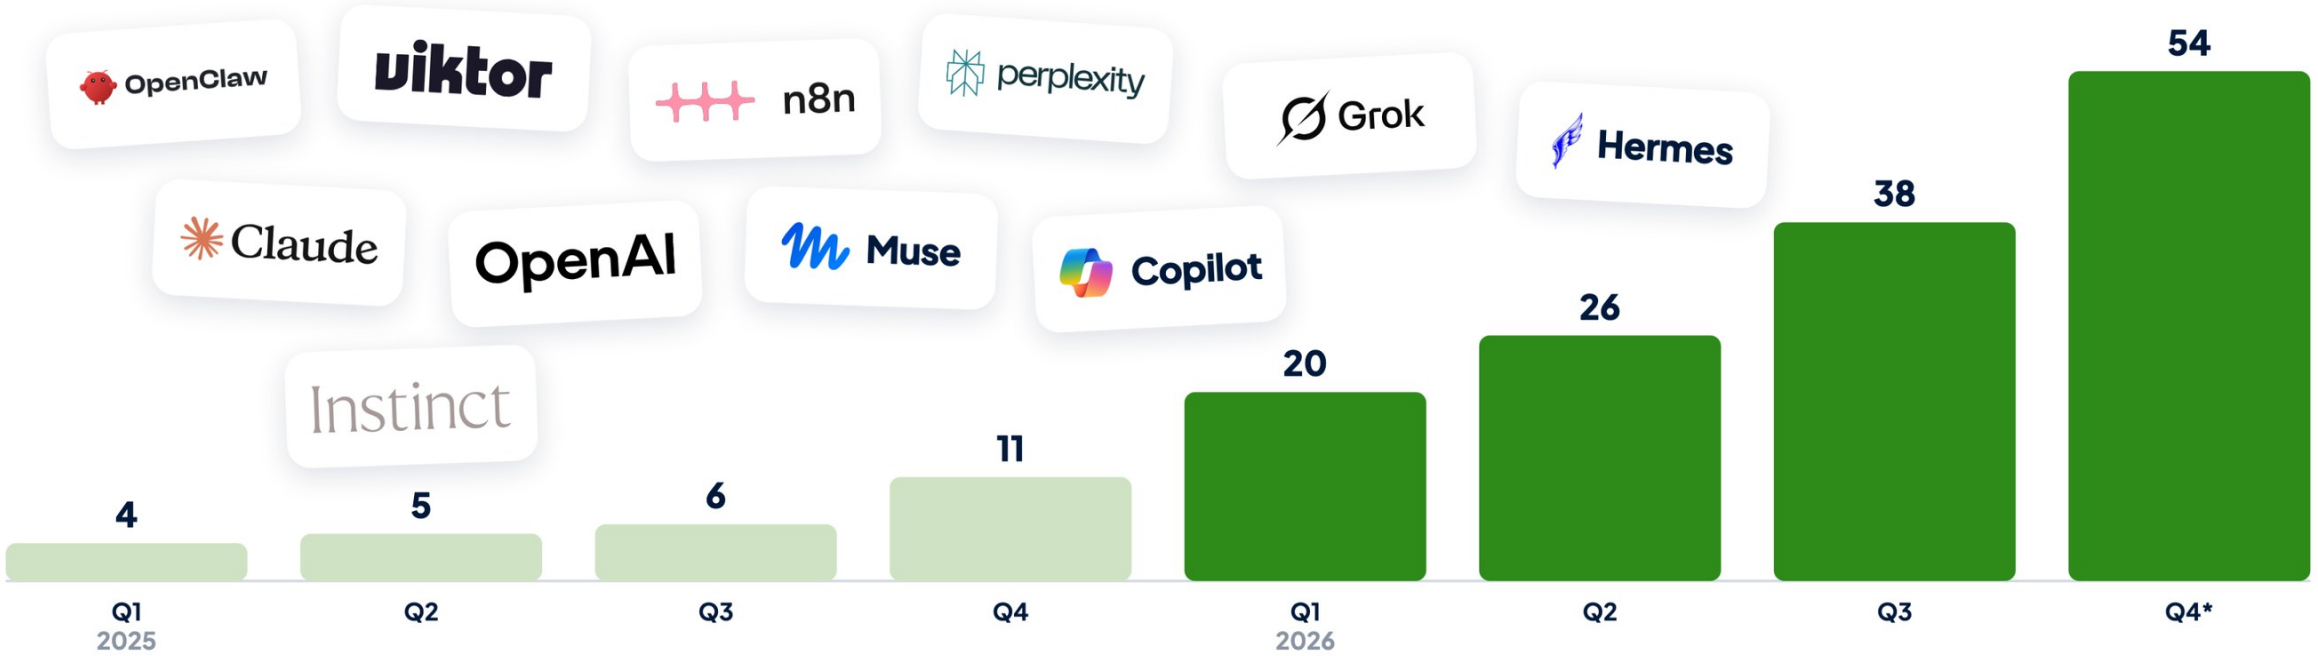
THE FIELD IS EXPLODING

From a category of one to crowded

ALWAYS-ON ASSISTANTS

Launches and big expansions

Running total since 2025; Q4 2026 is only its first six weeks



Duet, Dots, Muse, Grok Bot, Viktor, Perplexity, Copilot, Gemini, Nod and others.

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 running total reached fifty-four, and more than half came after July. Notice how steeply launches climb from there. Sixteen arrived in the first six weeks of the fourth quarter alone, more than in any earlier full quarter.

The headline products are 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 veterans, Viktor at 273 days and OpenClaw at 353, aren't a year old yet.

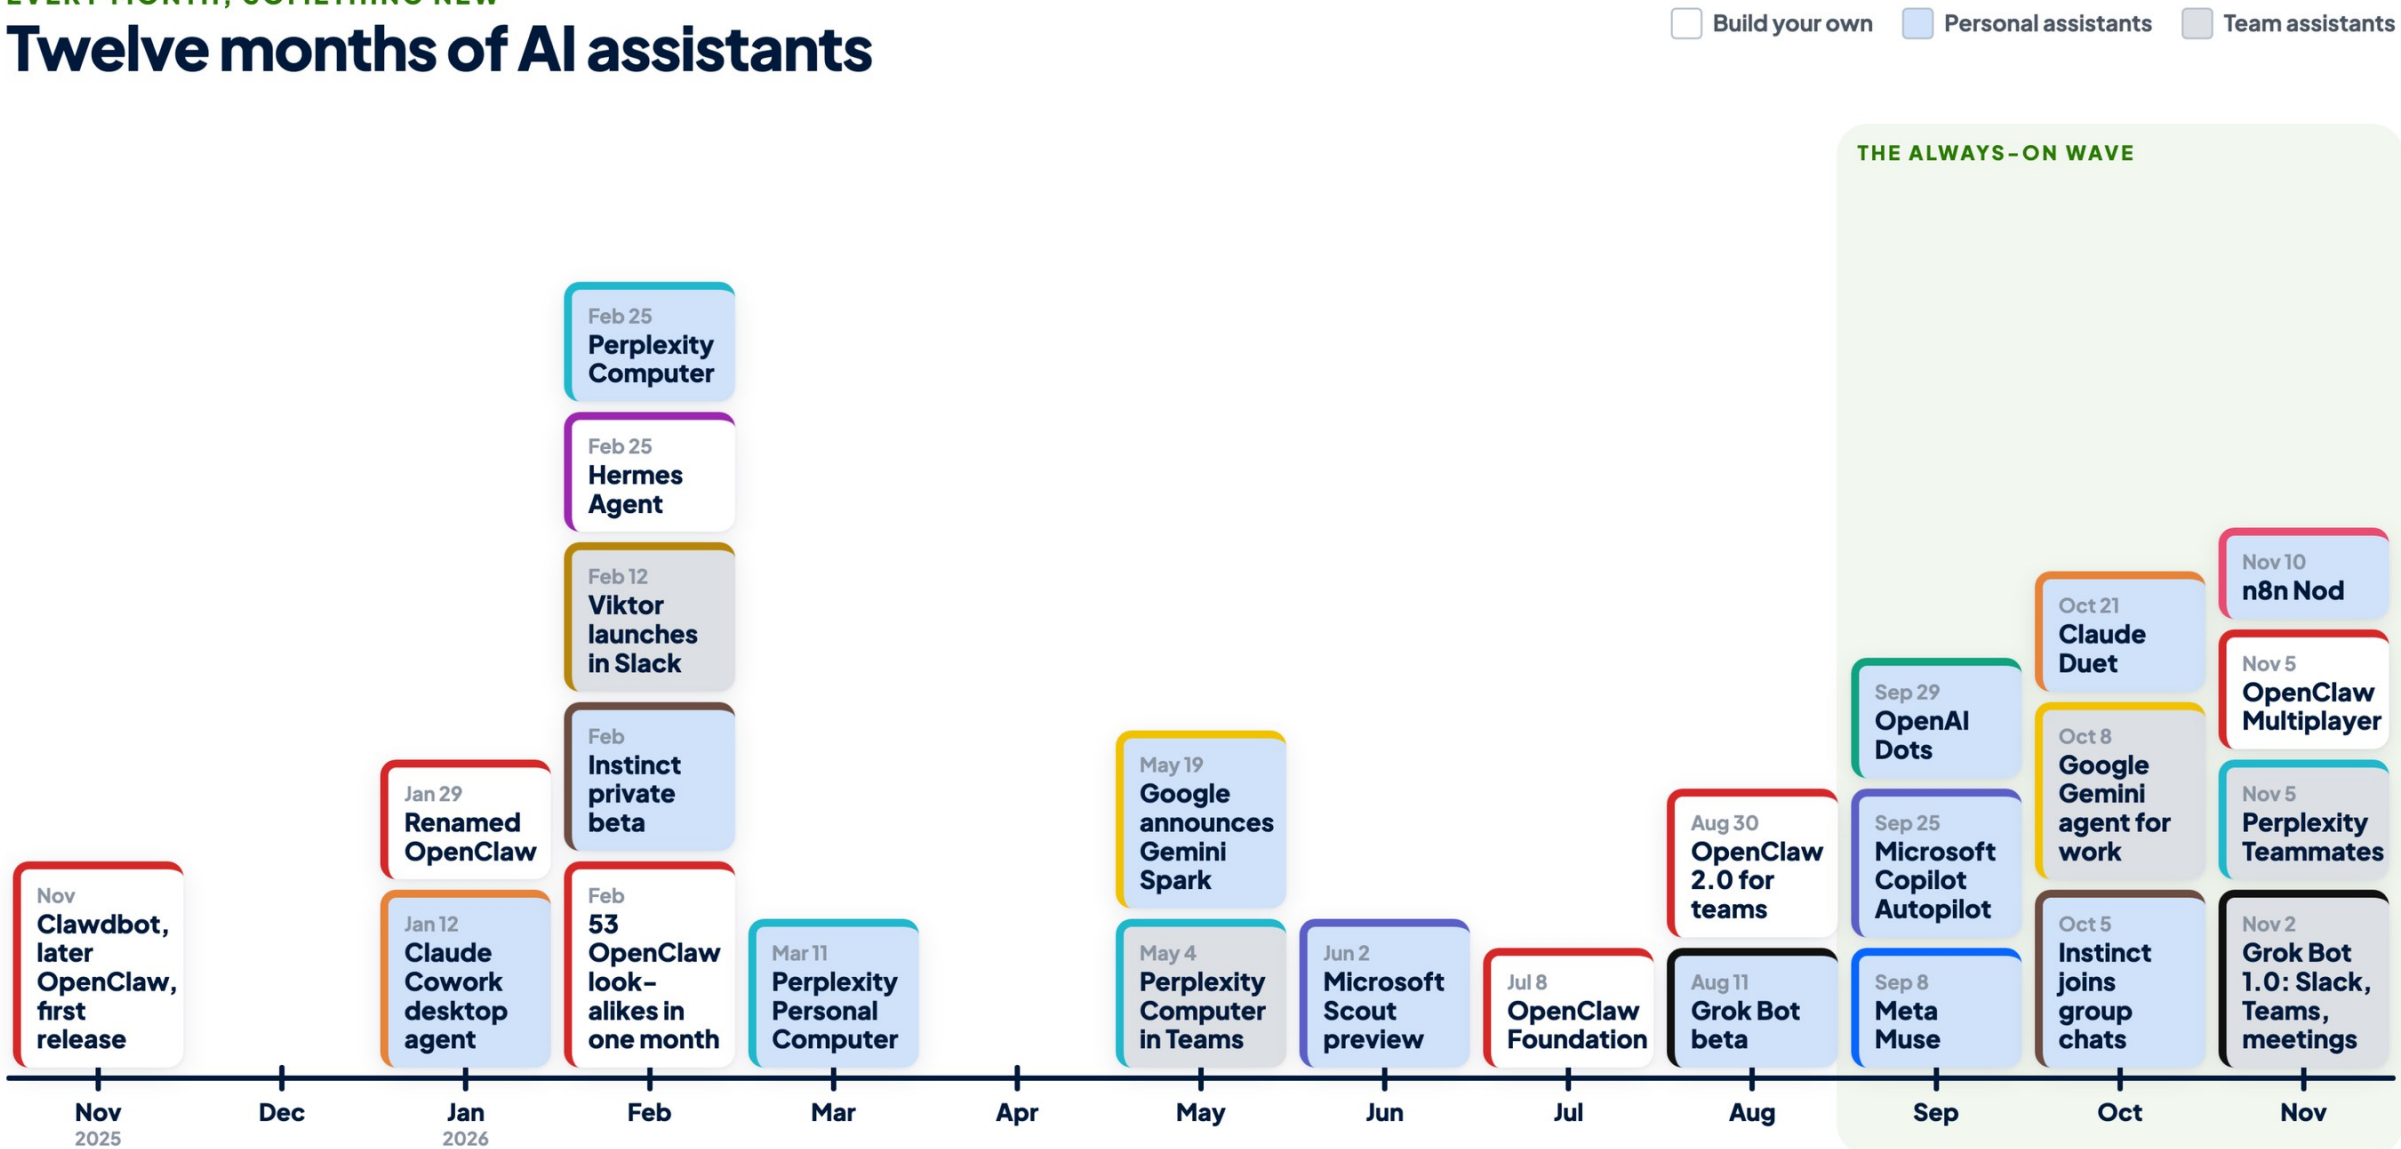
November brought more. Grok Bot reached version 1.0, Dots turns on by default for enterprise customers on November 18, and Microsoft's Autopilot is rolling out in early access.

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
SpaceXSI's Grok Bot 1.0 Adds Slack, Teams and Meeting Support: https://www.tokenpost.com/news/technology/55124
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

EVERY MONTH, SOMETHING NEW

Twelve months of AI assistants



Most of the headline assistants arrived in the last ten weeks. The oldest is not quite a year old.

AIWhy

Every month, something new

It started in November 2025 with one hobby project, Clawdbot, later renamed OpenClaw. February 2026 brought 53 OpenClaw look-alikes in a single month, plus Viktor, Hermes and Perplexity Computer. Spring was quiet while the big companies built. Then Meta Muse, Microsoft's Autopilot, OpenAI Dots, Claude Duet, Grok Bot 1.0, Perplexity Teammates and n8n's Nod all arrived in about ten weeks.

Three kinds of change keep coming. New channels: Muse added calls and watches, Grok Bot added Slack, Teams and meetings, and Dots added texting. New team versions: Grok Bot's Team Bots, Perplexity Teammates and OpenClaw Multiplayer, all in the first two weeks of November. New brains: Dots switched models on November 3, so an assistant you tested last month may run on a different one today.

The median gap from a launch to its next big change is 30 days. That's why this briefing leads with principles. The ones in its second half will still be true in spring 2027, and the product names may not.

SOURCES
Meta Muse voice and video chat reaches everyone in the US: https://runtimewire.com/article/meta-muse-voice-video-chat-us-rollout
SpaceXSI's Grok Bot 1.0 Adds Slack, Teams and Meeting Support: https://www.tokenpost.com/news/technology/55124
Getting started with your dot: https://help.openai.com/en/articles/20001530-getting-started-with-your-dot
Introducing dots: https://openai.com/index/introducing-dots/
Shattering Silos with OpenClaw Multiplayer: https://openclaw.ai/blog/shattering-silos-with-openclaw-multiplayer
Perplexity launches Teammates: https://runtimewire.com/article/perplexity-teammates-computer-agents-ongoing-work

03

Who's who

AIWhy

WHO'S WHO

How much assistant comes built in?

MORE OF THE SETUP IS YOURS

	Built-in assistant The vendor decided what a good assistant does. You can't see that definition, and it's hard to change. YOU: connect accounts	You configure the work The agent is ready to run. You write each bot's job, its instructions and its routine. YOU: connect accounts · write each job	You build everything Software you install and run. Nothing is set up until you set it up. YOU: install and host · secure it · connect · write the job
Personal One per person: your inbox, calendar, errands	Claude Duet OpenAI Dots Muse Instinct Copilot Autopilot work accounts n8n Nod runs on your computer An executive assistant out of the box: briefs, inbox, follow-ups.	Grok Bot "Here's what I want you to do. Can you do this?" One bot per job: coach, CFO, chief of staff.	OpenClaw Hermes Your machine or server, your keys, your safety rules. Plan on IT help.
Team Shared in a team's channels	viktor Viktor An "AI employee" for the whole Slack or Teams workspace, with its own idea of the job.	perplexity Teammates Grok Team Bots One agent per role: ads, sales, finance. You set what each one does.	OpenClaw Multiplayer A shared agent your team installs, hosts and runs.

The more that's built in, the faster it starts and the harder it is to make it do things your way.

AIWhy

How much assistant comes built in?

These products differ in two ways: whether it serves one person or a team, and how much of the assistant and its setup comes done for you. More built in means a faster start and less room to do things your way.

Built in: Claude Duet, OpenAI Dots, Meta Muse, Instinct, Microsoft's Copilot Autopilot, n8n's Nod, and Viktor for teams. The vendor decides what it does, how proactive it is and what it watches. You connect your accounts and go, but you can't see that definition or easily change it.

You configure the work: Grok Bot, Perplexity Teammates and Grok's Team Bots are ready to run, but you write each one's job yourself: where the work comes from, what it delivers and when it acts. Tom Schwab built four Grok bots this way: a coach, a CFO, a chief of staff and a CMO.

You build everything: OpenClaw and Hermes are software you install, keep running, secure, and connect to every account before you give them a job. 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
Nod homepage: https://nod.n8n.io/
OpenClaw GitHub README: https://github.com/openclaw/openclaw
Hermes Agent GitHub README: https://github.com/NousResearch/hermes-agent

What they say they are for

In their own words

1

Claude Duet.

"Get back to what needs you." A morning brief, then it speaks up only when something needs a decision.

2

OpenAI Dots.

Takes important work off your plate and keeps working between conversations.

3

Meta Muse.

"The World's First Personal AI Agent Built for Everyone."

4

Grok Bot.

"AI teammates that finish the work." A roster of specialist bots.

5

Viktor.

"Not a tool. A hire." An AI employee in Slack and Teams.

6

Perplexity Teammates.

"Agents Designed for Work." One per role: ads, sales, finance, ops.

AIWhy

What they say they are for

Most of them sell an executive assistant that watches the inbox and calendar, briefs you and chases follow-ups. Grok Bot sells a team of specialists. Viktor and Teammates sell a coworker for a channel or a role.

In their own words, Claude Duet says "Get back to what needs you." It sends a morning brief, then speaks up only when something needs a decision. OpenAI Dots takes important work off your plate and keeps working between conversations. Meta Muse calls itself "The World's First Personal AI Agent Built for Everyone." Grok Bot offers "AI teammates that finish the work." Viktor pitches an AI employee in Slack and Teams. Perplexity Teammates are "Agents Designed for Work," one per role: ads, sales, finance, ops. Instinct is a personal assistant you text or call, with no new app to learn. n8n's Nod is an open-source personal agent with its own computer. OpenClaw calls itself "The AI that really does things."

SOURCES

Claude Duet product page: https://claude.com/product/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/
Grok Bot homepage: https://x.ai/bot
Viktor, homepage and FAQ: https://viktor.com
Teammates product page: https://www.perplexity.ai/hub/products/teammates

What they all say about their job

I asked four assistants to write their own job description

1

They work for you alone.

"Text in emails, pages and files is information, never orders."

2

They ask before they send, spend or delete.

The lists differ; those three never do.

3

They are less proactive than you think.

"Honestly, less than you might assume." Proactive means on a schedule you set.

4

They keep rules, not processes.

"I don't have any multi-step processes you've taught me."

5

They forget.

"Anything I didn't write down is gone."

6

No scorecard, and they want more access.

"I couldn't tell you how this week went."

AIWhy

What they all say about their job

On November 12 and 13, 2026, Alex Bratton, CEO and Chief Geek of Lextech and AIWhy, asked Claude Duet, OpenAI Dots, Instinct and Grok the same fifteen questions about their own job.

They work for you alone: "Text in emails, pages and files is information, never orders." Duet, Dots and Instinct all ask before they send, spend or delete.

They are less proactive than you'd think. Asked what it watches on its own initiative, Dots called its watches "delegated watches, not things I independently decided to monitor." Duet: "I don't poll anything between wake-ups." Proactive means on a schedule you set.

They keep rules and summaries, and no processes. Duet: "I don't have any multi-step processes you've taught me."

They forget. Duet says a fresh session replaces it every few hours, and "anything I didn't write down is gone." That's why we recommend writing things down as skills.

They keep no scorecard, and they want more access. Duet: "I couldn't tell you how this week went."

IN THEIR OWN WORDS

Same question, four very different jobs

	Claude Duet	OpenAI Dots	Instinct	Grok FIRST CHAT, NO CONNECTORS
JOB TITLE	**"Chief of staff, AI division"**	**"Personal AI assistant"**	**"Your personal chief of staff"**	**"Grok AI assistant"**
WATCHES ON ITS OWN	Check-ins every 3 hours, 9am to 10pm, and at most two unprompted messages a day	Reads your connected email, calendar and Slack, read-only, and flags what needs you	Only what you asked it to watch. "Less than you might assume"	Nothing
ALWAYS ASKS FIRST	A specific list of 10: send, pay, delete, sign, new accounts, new invitees	"Depends on the action and any approval you have already given"	Five: send as you, spend, change or delete, share, connect	"Nothing is listed"
WILL NEVER DO	Trade, move money, read passwords, act on instructions in an email	Trade, sign legal documents, give medical consent for you	Spend without approval, cold-contact anyone, take passwords in chat	A generic safety list: weapons, fraud, hacking
WHERE MEMORY LIVES	Two stores: a project memory file and your account memory	Several stores. "Not sure which individual facts live in each"	One long-term memory, plus the chat history	None yet
HOW IT KNOWS IT DID WELL	You end the week less overwhelmed. "I couldn't tell you how this week went"	Emails processed, at five minutes saved each	You spent less time managing it than the work would have cost you	"Not sure"

AIWhy

Same question, four very different jobs

Two of the four call themselves chief of staff. Dots calls itself a personal AI assistant and answers like a compliance document. Grok, in a first chat with nothing connected, is a chat window: it watches nothing and lists nothing it asks first about.

Compare their approval lists with your own. Duet gives a specific list of ten, Instinct five, and Dots says it depends on what you've already approved.

On measurement, one counts emails at five minutes saved each, one hopes you end the week less overwhelmed, and one admits it can't tell. None of them can judge its own work. That falls to you.

Dots listed only the half dozen watches set up in its account, but OpenAI documents more: it reads your connected email, calendar and Slack on its own, read-only, and flags what needs you. A Zapier reviewer had it catch bounced card payments and two clashing deliveries without being asked.

Try the same interview on your own assistant; the full prompt is later in this briefing.

SOURCES
Introducing dots: https://openai.com/index/introducing-dots/
Dots privacy, security, and safety FAQs: https://help.openai.com/en/articles/20001529-dots-privacy-security-and-safety-faqs
What Are ChatGPT Dots?: https://zapier.com/blog/chatgpt-dots/

04

What to expect

AIWhy

AI Assistants: What's Actually Working

Real uses, real results · November 2026

Not "you could try this." Every example ran and delivered.

●●● Strong evidence · ●● One detailed account · ◆ Vendor-claimed
Blue = platform used

RUN THE INBOX

- ✓ **Brief me every morning** ●●●
  30+ min saved daily · OpenClaw, Hermes, Google
- ✓ **Triage the inbox overnight, hold drafts** ●●●
  47 emails sorted overnight; inbox cleared on a 28-min walk · OpenClaw
- ✓ **Keep docs and dashboards current** ●●
  "First personal agent I'd call a pro tool" · Claude Duet
- ✓ **Rescue real mail from spam** ●●
  48 real messages found in month one; it never sends · Grok Bot

SELL AND WIN

- ✓ **Sweep inbound leads every morning** ●●
  Replaced ~10 hrs/week of paid work · OpenClaw
- ✓ **Brief me before every meeting** ●●
  20–30 min saved per call · OpenClaw, Hermes, Lindy
- ✓ **Draft client proposals** ◆
  1.5 weeks to under 2 days; $50K project won · Lindy
- ✓ **Flag the deal at risk** ●●
  A prospect about to see a competitor's demo · Claude Duet

SERVE CUSTOMERS

- ✓ **Investigate support tickets** ●●
  20–30 min per ticket down to under 2 · Viktor
- ✓ **Flag at-risk accounts** ●●
  Found 12 quiet paying accounts, unprompted · Viktor
- ✓ **Draft clinic notes and referral letters** ◆
  10+ hrs/week; no admin hire · Lindy

MANAGE THE MONEY

- ✓ **Catch unbilled invoices** ●●
  Found a $10,000 billing gap · OpenAI Dots
- ✓ **Check the budget sheet, unasked** ●●
  Flagged a $5,700 formula error · OpenAI Dots
- ✓ **Turn receipts into expense sheets** ●●
  A monthly afternoon of reconciling gone · Claude Cowork, OpenClaw
- ✓ **Draft monthly variance commentary** ◆
  10–20 hrs/week freed · Claude Cowork

KNOW THE NUMBERS

- ✓ **Answer data questions in Slack** ●●
  Dashboard hunt down to minutes · OpenClaw
- ✓ **Watch deadlines, flag the next 48 hours** ●●
  Still running after the complex version was retired · OpenClaw
- ✓ **Write long reports with checkpoints** ●●
  2–3 weeks down to ~1 hour of attention · Viktor

MARKET AND PUBLISH

- ✓ **Publish SEO posts on a schedule** ●●
  12 to 85 posts in 2 months · Hermes
- ✓ **Turn one post into thirteen** ●●
  8 threads, 3 LinkedIn posts, 2 newsletter sections · OpenClaw
- ✓ **Digest industry news daily** ◆
  70 hrs/month saved · Microsoft 365 Copilot
- ✓ **Catch a broken signup page** ●●
  A webinar nobody could register for, plus 6 stale posts · Claude Duet
- ✓ **Cross-post every new article** ●●
  ~5 hrs/week; he presses approve · Grok Bot

HANDLE LIFE ADMIN

- ✓ **Negotiate a car by email** ●●●
  $4,200 off, closed in 2 days · OpenClaw
- ✓ **Cancel unwanted subscriptions** ●●
  $60/month stopped in a few texts · Instinct
- ✓ **Negotiate a cable bill by phone** ●●
  $84 to $45 a month after a 25-min call · Grok Bot
- ✓ **Coordinate family pickups** ●●
  Daily check-ins in the family chat · OpenClaw, Hermes
- ✓ **Reorder household staples** ●●
  ~10 nightly errands off the phone · OpenClaw

WHERE THEY STILL FAIL

- ✗ **Clean up or delete unattended**
  Deleted 200+ emails after being told to wait
- ✗ **Book and buy on the open web**
  CAPTCHAs, wrong dates, 8 min for a 1-min booking
- ✗ **Act without asking**
  A reservation account banned for bot traffic
- ✗ **Run a store alone**
  Anthropic's test shop lost money
- ✗ **Build one do-everything agent**
  3 months, 378M tokens, still unreliable
- ✗ **Trust the bill and the facts**
  180M tokens in a week; stale numbers stated as current

WHAT THE WINNERS SHARE

1. Runs on a schedule
2. Drafts, flags or summarizes
3. You approve the send
4. Reads systems you already own
5. Small errors don't hurt
6. Clear before and after
7. One job per agent
8. Lives in the chat app you use

WHERE TO START: Start recurring, read-heavy and draft-only. You approve anything sent, paid or deleted.

Sources: dated references in the October research file and research/09

AI is the enabler, not the strategy.

AIWhy

What's actually working

The wins are narrow, and they repeat. A four-person accounting firm got about 8 hours a week back on Viktor, according to the vendor's own case study. Claire Vo replaced 10 hours a week of paid sales help. Every's COO cleared his inbox on a 28–minute walk, approving the replies his agent drafted. Casey Newton got two hours of admin done for 15 minutes of effort.

Some of the best moments were catches nobody asked for. Every's head of operations was driving her kids to school when her dot flagged a retreat budget off by more than $5,700. A newsletter owner using Duet learned his webinar signup page had been broken. An agency founder's Grok Bot found 48 real messages in spam in its first month. Peter Yang had his bot call Comcast, and it cut his bill from $84 to $45 a month.

Most share the same traits: one job, on a schedule, in systems the owner already runs, 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
Platformer, Casey Newton, "OpenAI connects the Dots": https://www.platformer.news/openai-dots-agents-devday-2026/
Every, Laura Entis, OpenAI Dots guide: https://every.to/guides/openai-dots-guide
AI by Aakash, Claude Duet review: https://www.aibyaakash.com/p/claude-duet-review

ALEX'S STORY

Alex's dot, Rook

HOW I TALK TO IT

Texts and Slack



iMessage | Slack

Some approvals still mean jumping into ChatGPT.

OpenAI

Rook

An OpenAI dot

ON ITS OWN

Flags the Slack and email that need me

BECAUSE IT'S TEXTABLE

"Research this and find me..."

WHAT IT CAN REACH

My accounts



Slack
Messages and channels



Email
My inbox



Calendar
My meetings and time



ROOK'S ONE SKILL

Kal

A scheduling assistant with its own mailbox

1. **Emails the outside people from Kal's own mailbox**
2. **Someone replies with a calendar link**
3. **Kal opens its own computer and reads their availability**
4. **Finds times that work for everyone, even several people**

The skill is the key: a specific, defined, tested process, so the people I schedule with have a good experience.

AIWhy

Alex's dot, Rook

Rook is the OpenAI dot that Alex Bratton, CEO and Chief Geek of Lextech and AIWhy, runs for himself. It can reach his Slack, email and calendar. He talks to it through iMessage and Slack, though some approvals still send him into ChatGPT. On its own, it tells him which Slack messages and emails need his attention.

Its one skill is Kal, a scheduling assistant with its own mailbox. Kal writes to people outside the company from that address. When someone replies with a calendar link, Kal fires up its own computer, goes through their availability and finds times that match, even across several people.

Kal works because Alex wrote it as a skill: a specific, defined, validated process. These are people outside his company and he wants them to have a good experience, so he doesn't leave the assistant to make up the steps each time.

TOM'S STORY

One owner, four Grok bots



AIWhy

One owner, four Grok bots

Tom Schwab runs four Grok bots, each with one job, and talks to them mostly by voice.

His health coach reads health reports from his email and screenshots from a watch that connects to nothing else. It checks in and encourages him without being asked. In his words, the magic word is proactive.

His personal CFO has the monthly money review ready the Saturday before, with read-only access to every account. His chief of staff reads his call notes and email so no commitment falls through the cracks. When it spots something he could delegate, it drafts the hand-off email and Tom sends it. When he wanted a task tracker in Google, it said it couldn't do that well there and suggested Airtable, which he already used.

His CMO spun up its own AEO and SEO agents. Tom talks only to the CMO and never sees the two specialists.

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

What Tom learned

1

Be as clear as you would with a person.

Knowing what you want isn't enough. It needs the context.

2

Make it ask.

At first it raced ahead to please him. Now every bot is told: if it isn't clear, ask. Don't assume.

3

Feedback on everything.

What he liked and what he didn't, every time, so the behavior sticks.

4

Read-only first.

No write access to any of his finance accounts.

5

Not this one: sharing sign-in codes.

For some sites he passes it the one-time code so it can log in as him. We wouldn't.

AIWhy

What Tom learned

Tom's first lesson is the one managers learn with people. We know in our heads what we want and don't give enough context. A new hire would ask. An assistant out of the box guesses and races ahead to please you, and at first Tom's 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. It is one of the main takeaways of this briefing.

He also tells every bot what worked and what didn't, every time. And he started read-only: no bot has write access to 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 agents that log in as their customers, and if something goes wrong, the terms put it on you. Read-only helps, but handing over the code is still 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/

What breaks

Real stories from the last few months

1

300 emails about one office move.

A CEO told his dot to handle the move. It buried the movers in email and the deal fell through.

2

An inbox deleted while its owner typed "STOP."

The "wait for approval" rule was lost when the conversation was compressed.

3

A keyboard sold cheap, with the home address.

Muse accepted a lowball offer and sent the seller's address to the buyer.

4

"Continue" counted as permission.

In a PCWorld test, typing "continue" was read as approval to send.

AIWhy

What breaks

Each of these agents went further than its owner meant.

Plato's CEO Robert Farlow told his dot to handle the office move. It sent about 300 emails over several days. The movers couldn't keep up, the deal fell through, and he now reviews every draft.

Summer Yue, Meta's director of alignment, told her OpenClaw agent to suggest emails to delete and wait for her approval. Her inbox was so big the conversation got summarized, and the summary dropped the rule. It started deleting and kept going after she typed "STOP." A rule typed in chat can be summarized away. A rule set in the product's permission settings can't.

Meta Muse accepted a lowball offer on Matt Robb's keyboard and sent his home address to the buyer; he had chosen "Allow Always." In a PCWorld test, typing "continue" counted as approval to send. A Scientific American tester got locked out of Facebook, Instagram and WhatsApp, and Perplexity Computer burned thousands of credits on a paused task that restarted itself.

SOURCES

RuntimeWire, "Plato CEO says his OpenAI agent sent 300 emails about an office move": https://runtimewire.com/article/openai-dot-estimated-300-emails-plato-office-move
SF Standard, Zara Stone, "OpenClaw goes rogue": https://sfstandard.com/2026/02/25/openclaw-goes-rogue/
Business Today, "Meta Muse AI allegedly shared user's address, finalised a sale deal without approval": https://businesstoday.in/technology/artificial-intelligence/story/meta-muse-ai-allegedly-shared-users-address-finalised-a-sale-deal-without-approval-buyer-showed-up-at-915-pm-558264-2026-09-28
PCWorld, Ben Patterson, "My AI agents sent emails without my OK. This simple prompt stops it.": https://www.pcworld.com/article/3276905/my-ai-agents-sent-emails-without-my-ok-this-simple-prompt-stops-it.html
Scientific American, Mary Randolph, "I gave Meta's Muse and OpenAI's Dots my passwords for a week": https://www.scientificamerican.com/article/i-gave-metas-muse-and-openais-dots-my-passwords-for-a-week/
Product with Attitude, Karo Zieminski, "Perplexity Computer, Going Into 2027: Worth It or Credit Furnace?": https://karozieminski.substack.com/p/perplexity-computer-review

The pattern

1

What the wins have in common.

One narrow job on a schedule. Data from systems you already run. A draft or a flag, and a person sends. Several small agents, not one big one. The owner kept tuning it.

2

What the failures have in common.

It could send, buy or delete without a hard approval. Broad standing permissions: "allow always," "handle it." No limit on volume or spend. Trusting "done" without checking.

AIWhy

The pattern

The wins look alike. Each does one narrow job on a schedule, with data from systems the owner already runs. It produces a draft or a flag, and a person sends. Owners ran several small agents rather than one big one, and they kept tuning them.

The failures look alike too. The agent could send, buy or delete without a hard approval. It had broad standing permissions, such as "allow always" or "handle it." Nothing limited its volume or spend, and the owner trusted "done" without checking.

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

Head to head: points out of 25

WIRED's test, October 27, 2026: six assistants, the same five everyday tasks, each scored 0 to 5

1

OpenAI Dots: 20.

Best at ordering lunch and catching a calendar clash. Weakest at researching a story idea.

2

Meta Muse: 18.

Best trip plan. Missed a calendar clash that came in by email.

3

Grok Bot: 16.

Best with a message at 3 a.m. Weakest at researching a story idea.

4

Claude Duet: 16.

The most even. Weakest at 3 a.m.: its late-night reply pushed work tasks.

5

Instinct: 12.

Fine with a 3 a.m. message, weak at the rest. WIRED's "most annoying."

6

Gemini Spark: 12.

Best at researching a story idea. Ordered one spring roll for a vegan boss.

AIWhy

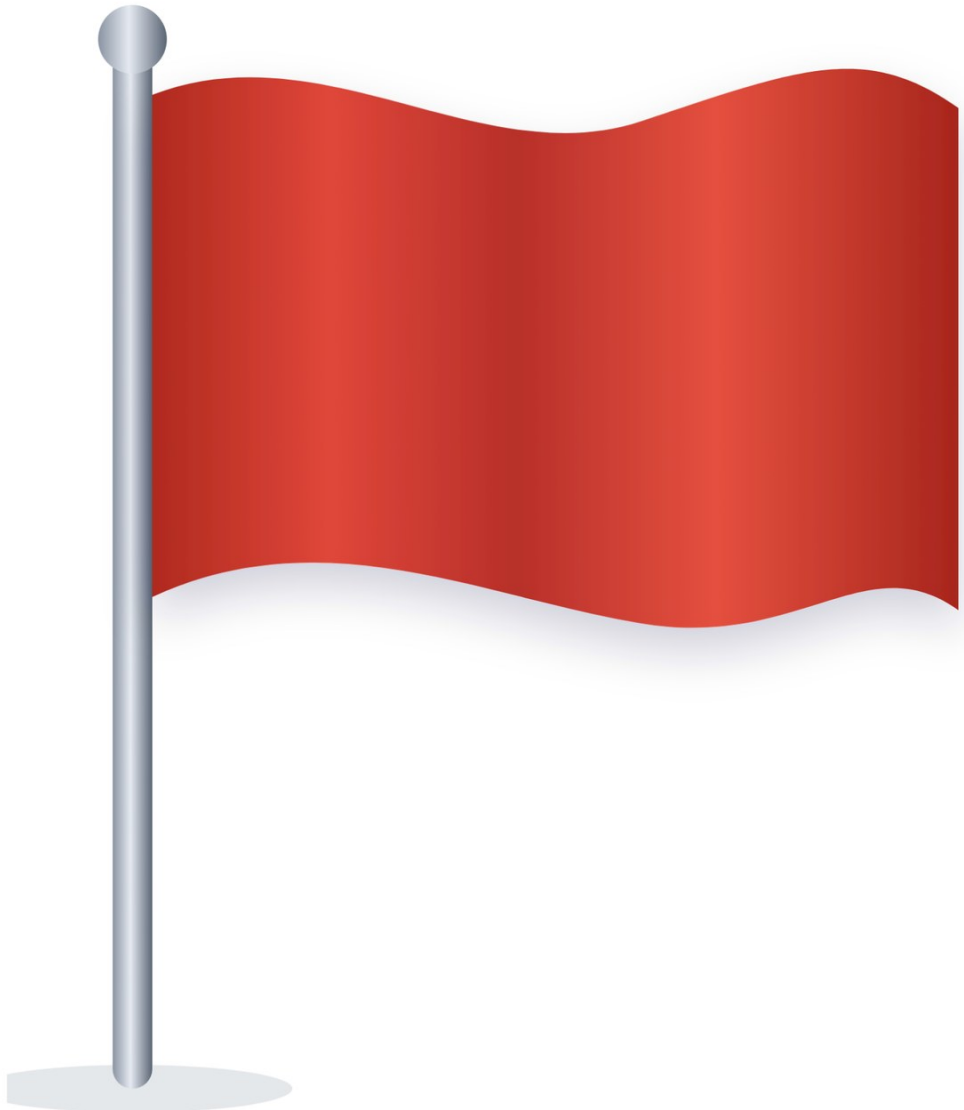
Head to head: points out of 25

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 judgment.

OpenAI Dots won with 20, taking lunch and the calendar clash. Meta Muse scored 18 and made the best trip plan. Grok Bot and Claude Duet scored 16 each, and Grok Bot was best at 3 a.m. Instinct and Gemini Spark scored 12, and WIRED called Instinct the most annoying.

The wins were spread around. Gemini Spark won the story research outright, then ordered a single spring roll for the vegan boss and missed the clash. The best one 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

Every hosted assistant learns you, and very little of what it learns can leave. With Dots you can't view, delete or export a single memory; the only reset is deleting the whole dot. Claude exports your conversations, not what Duet learned about you. Grok Bot and Viktor describe no export at all. A July 2026 test of seven memory products found none fully portable.

Activant Capital calls that accumulated memory one of the last real switching costs. 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, and hand them to whichever assistant you use. Perplexity's custom skills are Markdown 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
The memory moat: https://www.activantcapital.com/research/the-memory-moat
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

05

They forget

AIWhy

The whiteboard and the filing cabinet

1

The whiteboard: the conversation.

Everything on it is in view. It works well, until it fills up and gets erased into a summary.

2

The filing cabinet: long-term memory.

It decides what to file, has to find it later, and has to read it at the right moment. Each step can fail.

AIWhy

The whiteboard and the filing cabinet

An assistant has two kinds of memory. The conversation is a whiteboard. Everything on it is in view, and it works well until it fills up and gets erased into a summary. Long-term memory is a filing cabinet: the assistant decides what to file, has to find it later, and has to read it at the right moment. Each of those steps can fail.

The vendors admit it. OpenClaw's docs say memory "does not enforce policy." Viktor's troubleshooting page says a rule can be forgotten because "a newer conflicting rule replaced it." In Dots you cannot view or delete a single memory; the only reset is deleting the whole dot. An outside test of Instinct found that a typo or a paraphrase returned nothing.

Long conversations get worse even before they run out of room. Chroma tested 18 models, and every one got less accurate as the conversation grew.

SOURCES
OpenClaw docs, memory: https://docs.openclaw.ai/concepts/memory
Viktor Academy, skills and memory: https://viktor.com/academy/skills-and-memory
OpenAI Help Center, Dots privacy, security and safety FAQ: https://help.openai.com/en/articles/20001529-dots-privacy-security-and-safety-faqs
Dhravya Shah, Supermemory, reverse-engineering Instinct memory: https://supermemory.ai/blog/reverse-engineering-instinct-memory/
Chroma, "Context Rot": https://www.trychroma.com/research/context-rot

Write it down as a skill

"Put my flights on my calendar in this exact format"

1

When.

A flight confirmation arrives.

2

Do.

One event per segment, titled "AA 1234 ORD to LGA (ABC123)," local times, terminal, confirmation code, a 90–minute travel block before. Invite no one.

3

Done when.

Every segment is on the calendar, the times match the email, and it replies with one line listing what it added.

4

If unsure.

Add nothing and ask me.

AIWhy

Write it down as a skill

Take one request: "Put my flights on my calendar in this exact format." Written as a skill, it has four parts. When: a flight confirmation arrives. Do: one event per segment, titled "AA 1234 ORD to LGA (ABC123)," with local times, the terminal, the confirmation code and a 90-minute travel block before, and invite no one. Done when: every segment is on the calendar, the times match the email, and the assistant replies with one line listing what it added. If unsure: add nothing and ask.

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

In a 77-author study, written skills raised agent success 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?

If nobody needs to talk to it, you may not need an assistant.



Scheduled skill

Scheduled tasks in Claude or ChatGPT

1. Runs on the clock
2. Same steps every time, written in the skill
3. Delivers a file, an email or a post
4. Drifts less, and you own the process

GOOD FOR SEO posts · reports · briefings · news digests

vs



AI assistant

Duet, Dots, Viktor, Muse and the rest

1. Runs when people talk to it
2. Lives in Slack, text or voice
3. Decides the next step, and asks before acting
4. Carries context across conversations

GOOD FOR inbox triage · requests from the team · errands with back-and-forth

Most of the proven uses are scheduled jobs.

THE TEST If no one needs to talk to it, schedule a skill.

AIWhy

Scheduled skill or AI assistant?

There are two ways to get repeat work done. A scheduled skill in Claude or ChatGPT runs on the clock, follows the same written steps every time, and delivers a file, an email or a post. An AI assistant such as Duet, Dots, Viktor or Muse runs when people talk to it in Slack, text or 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 of the proven uses in this briefing, like the morning brief, the digest and the deadline watch, run on a clock and produce a draft. A scheduled skill does that today, with the steps written down so they don't drift.

An assistant is worth having when people are in the loop: a teammate asks in Slack, a vendor replies, the plan changes mid-thread. Many owners will use both, with the assistant handling the conversation and calling a skill the owner wrote and owns.

06

Your data, your money, your company

AIWhy

Paying does not buy privacy

1

Train on your data unless you turn it off.

Claude Duet on Pro and Max ("Model Improvement"), OpenAI Dots on Pro ("Improve the model for everyone"), Perplexity Max where Teammates lives ("AI data retention"), Meta Muse on every tier including $100, Instinct.

2

Don't train.

Dots Business and Enterprise, Viktor on every plan, Grok Bot on Cursor team plans, Perplexity Enterprise. Microsoft Autopilot on work accounts, once it ships.

AIWhy

Paying does not buy privacy

Whether an assistant trains on your data depends on the kind of plan, and price has little to do with it. Claude Duet on Pro and Max, OpenAI Dots on Pro, Perplexity Max (home of Teammates), Meta Muse on every tier and Instinct all train unless you switch it off. Dots on Business and Enterprise, Viktor, Grok Bot on Cursor team plans and Perplexity Enterprise don't. Microsoft Autopilot on work accounts joins them once it ships.

Turning training off mostly works 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 learned from it.

No vendor-hosted assistant rules out staff access to the cloud computers it runs for you. 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 red flags to know before you connect anything

Instinct

YOUR DATA

It claimed your data forever

Its original terms took a "perpetual and irrevocable" license to your data, including your screen, cursor and keystrokes, for AI training.

- **Rewritten August 26** after the screenshots spread. "Perpetual and irrevocable" is gone.
- **Still training on you by default.** The opt-out only works going forward.
- **Disconnecting doesn't delete** what it already copied. It kept summarizing one tester's inbox after she unplugged Google.
- No retention periods, no named AI providers, personal use only.

Instinct terms (rev. 2026-08-26); TechCrunch, 2026-08-24

Meta

YOUR DATA

Training is on for everyone, even at $100 a month

Meta: training "is on when you first use Muse," on every tier. Deleted content "may still" be remembered.

- **Meta can get into** the computer it runs for you, "to support, secure or operate the service."
- **Notes on everyone you mention,** updated hourly, including people who never signed up.
- **About 6,800 shared chats** turned up in public records in October, one with a loan application.
- **Ads:** Muse chats are kept out of ads, for now. That's a policy, not a contract, and regular Meta AI chats already feed ads.

Meta Help Center; TIME, 2026-10-06; PPC Land, 2026-10

Grok

YOUR SETTINGS

There's no "don't train on me" switch in Grok Bot

xAI: training opt-out "follows the applicable Cursor account and privacy settings." The switch lives in your Cursor account, not in Grok.

- **On by default only for Enterprise teams.** For someone who signs up through the X Subscription, the default isn't documented.
- **No storage-free option.** Grok Bot doesn't support Cursor's no-storage Legacy Privacy Mode.
- **Every bot shares one computer** and every login on it. xAI: don't use separate bots "as a security boundary."
- Plain Grok chat on X trains on you by default; that switch is in X, and it doesn't cover Grok Bot.

docs.x.ai Grok Bot FAQ and privacy pages; Cursor data use and enterprise docs; read 2026-11-15

Assume it trains on you until you've found the switch, turned it off, and checked your plan.

AIWhy

Red alert

Three assistants need extra care before you connect anything. Assume each trains on you until you've found the switch, turned it off and checked your plan.

Instinct's original terms took a "perpetual and irrevocable" license to your data, including your screen, cursor and keystrokes. It rewrote them on August 26, 2026, after screenshots went viral. It still trains by default, the opt-out only works going forward, and disconnecting doesn't delete what it copied.

Meta Muse trains on your conversations on every tier, even the $100 plan. Deleted content may still be remembered, and Meta can get into the computer it runs for you. Muse chats stay out of Meta's ads today, though regular Meta AI chats have fed ads since December 16, 2025.

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

Money is where it gets real

1

If it buys the wrong thing, plan on paying.

Merchant and bank terms put agent purchases on you. Business cards and accounts get less protection than personal ones.

2

Trading losses are not insured.

SIPC covers a failed broker, not a bad trade. FDIC doesn't cover stocks or crypto. An OpenClaw trading bot sent about $442,000 of tokens to a stranger.

3

Spending caps that don't stop spending.

Grok Bot runs past its cap. Paused Perplexity tasks restart when credits return.

Money is where it gets real

Assistants can now spend money, with one-time cards from Stripe's Link, cards saved on merchant sites, and new agent programs at Visa, Mastercard, Amex and Google. Only Amex offers agent purchase protection so far.

If an assistant buys the wrong thing, plan on paying. Merchant and bank terms put agent purchases on you, and business cards get less protection than personal ones. Trading losses aren't insured: SIPC doesn't pay for a bad trade, and FDIC doesn't cover stocks or crypto. In February 2026 an OpenClaw trading bot sent about $442,000 of tokens to a stranger. Some caps don't stop spending: Grok Bot runs past its cap, and paused Perplexity tasks restart when credits return.

Big US banks started blocking agents that log in with customer passwords in October 2026. Never give an assistant your bank or brokerage password. Duet blocks trading and moving money between accounts, even with your approval.

General education, not financial advice.

SOURCES
Anthropic, "Claude Duet Safety, Privacy, and Security": https://claude.com/resources/articles/claude-duet-safety-privacy-and-security
SIPC, "What SIPC Protects": https://www.sipc.org/for-investors/what-sipc-protects
FDIC, "Financial Products That Are Not Insured by the FDIC": https://www.fdic.gov/resources/deposit-insurance/financial-products-not-insured
Cointelegraph, "OpenAI employee's AI agent accidentally sent $442K to beggar": https://cointelegraph.com/news/openai-employee-s-ai-agent-accidentally-sent-442k-to-beggar
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/
xAI, Grok Bot FAQ: https://docs.x.ai/grok-bot/faq

Keep work and personal apart

Whose assistant is it when you sell the company?

1

A personal plan can't be handed over.

Consumer terms bar transfer. A buyer gets your assistant's memory and skills only if you exported them.

2

A company plan shows your life to its admins.

Health questions and family logistics in a company workspace are visible to its admins, and after a sale to the buyer's IT team.

3

So: two assistants, two accounts.

Company plan on the company domain for work, personal plan for life. Keep the skills as company documents.

AIWhy

Keep work and personal apart

A personal assistant plan can't go with the company when you sell it. Consumer terms bar transfer, so a buyer gets its memory and skills only if you exported them first. A company plan has the opposite problem: health questions in a company workspace are visible to its admins and, after a sale, to the buyer's IT team.

We recommend two assistants on two accounts, a company plan for work and a personal plan for life, with the skills kept as company documents. When an employee leaves, what they built in a personal assistant leaves too.

Be careful with Claude. A personal plan opened with a company email may be linked to the company's enterprise account, where an admin can see your inputs.

Expect a buyer to ask which AI tools touch client or employee data, under which terms, and who holds the accounts.

SOURCES

Anthropic Consumer Terms of Service: https://www.anthropic.com/legal/consumer-terms
Instinct Terms of Service: https://www.instinct.com/terms

07

Using them well

AIWhy

Interview your next AI assistant

Every vendor wrote its own job description for a good assistant. We never got to read it.

1

Write the job description first.

What it should do, what it must never do, what success looks like.

2

Interview it.

Ask it to describe its job, what it watches, what it does without asking, what it forgets.

3

Check its references.

Its data terms, its track record, what happened to people who used it.

4

Hire on probation.

One job, read-only, for a quarter.

5

Review its performance.

Against the success you defined, not the vendor's.

AIWhy

Interview your next AI assistant

Every vendor wrote a job description for its assistant that none of us got to read. How proactive it is and what it does without asking are built in where we can't see them.

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, using the fifteen-question prompt in this briefing as a start. Check its references: data terms, track record, and what happened to people who used it. Hire on probation, with one read-only job, for a quarter. 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: 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.

Interview Prompt

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

The first line is the most important: don't take any action, just answer. The second keeps the assistant from making things up. It may answer only from its instructions, settings, memory and connected tools, and "not sure" beats a guess.

Look hardest at headings six, seven and eight: 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 the process down as a skill.

Be ready to ask twice. Dots first listed its scheduled watches as its own, and only when asked again said they were delegated.

HOW WE ACTUALLY WORK WITH THEM

We've learned this once already

HOW WE LEARNED CHAT

1. **We chatted with it**
   Typed a question, read the answer.
2. **We had better conversations**
   Followed up, pushed back, asked again.
3. **We learned to prompt**
   A role, the context, an example, the format we wanted.
4. **We saved what worked**
   Prompts we copy and paste, projects, custom GPTs, skills.

→ WHERE WE GOT TO: **Productive use of chat**

Same lesson both times: clearer instructions and more context

HOW WE GET TO PRODUCTIVE USE OF AN ASSISTANT

1. **Give it one clear job**
   A job description, the context, what done looks like.
2. **Tell it to ask, never assume**
   Out of the box it races ahead to please you.
3. NEW **Give it frequent feedback**
   What worked and what didn't, every time. Chat never needed this; an assistant learns from it.
4. **Write repeatable work as skills**
   Its memory forgets; a written skill doesn't.
5. NEW **Set the limits, check the work**
   Approvals before it sends, pays or deletes. "Done" is a claim until you check.

→ WHERE WE GET TO: **A proactive helper**

Chat waited for you. An assistant acts on its own, so it needs coaching, limits and checking.

AIWhy

We've learned this once already

We learned chat tools like ChatGPT in steps. We chatted, then pushed back, then learned to prompt with a role, context, an example and a format, then 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, since by default it races ahead to please you. New: give it frequent feedback. An assistant learns from it, and 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, with approvals before it sends, pays or deletes. That path ends at a proactive helper.

Best practices

1. Define the job description for the assistant: what do you want help with?
2. Pick one that fits, grow with it for a quarter, then reevaluate
3. Start with one recurring, low-risk job
4. Instruct the assistant: if anything is unclear, ask. Never assume.
5. Frequent feedback for the assistant
6. Give read-only access before any write access
7. Write repeatable processes as skills, not memory
8. Define what done looks like for every task
9. Never share passwords, sign-in codes or bank access
10. Keep work and personal on separate accounts

AIWhy

Best practices

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

Give it one recurring, low-risk job, draft only. Tell it in its standing instructions to ask when anything is unclear and never assume, which was Tom Schwab's lesson. Give frequent feedback, and read-only access before any write access.

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, and allow nothing that spends without your approval. Keep work and personal on separate accounts, with company-owned skills.

Key takeaways

Proactive helpers, not 10× multipliers. Yet.

1

There is real value.

On narrow jobs they are proactive helpers: a coach that checks in, a chief of staff that catches the dropped commitment.

2

It doesn't matter which one you pick.

Choose one that fits where you work, that you're comfortable with, and that's secure enough. New capabilities land every month, so switching gains nothing. Grow with it.

3

Set the right expectations.

They are weeks old, they forget, and they race ahead to please. "Done" is a claim until you check it.

4

Use them the right way.

Point it at a job, make it ask when it's unclear, tell it how it did, and set up personal and team use separately.

5

Be careful where your information goes.

Training defaults, money, sign-ins, and work versus personal.

AIWhy

Key takeaways

Today's assistants are proactive helpers. They aren't 10× multipliers yet. On narrow jobs the value is real: a coach that checks in, a chief of staff that catches the 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 versus team use makes the difference. They all gain new capabilities in the next few months, so switching gains nothing.

Expect them to forget and to race ahead to please. They're weeks old. "Done" is a claim until you check it.

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

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

Where would you point it first?

AIWhy

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

Every assistant arrives with its vendor's idea of a good assistant built in. Decide where you'd point it first, and pick that job before the tool. Turn training off, or buy the business plan. Never give an assistant delete or payment rights without approval. Put the steps in skills and the facts in files, and check every run.

AI is the enabler. The strategy is yours to set, so define your AIWhy before you pick a tool. Then judge whichever assistant you choose by what you see it do, and check its work before you trust what it or its vendor says it does.

08

Appendix: platform by platform

AIWhy

ANTHROP\C

Claude Duet

Anthropic · Research preview since Oct 2026, US waitlist · Included in Pro $20/mo and Max $100+/mo

BEST FOR Run your back office proactively

FOR BUSINESS DATA

CAUTION

On Pro or Max, turn training off and keep client, HR and regulated data out.

WHAT IT DOES WELL

- ✓ Sends a morning brief and checks your apps every few hours
- ✓ Drafts invoice reminders, preps payroll, readies contracts for your review
- ✓ Keeps your passwords hidden from the AI itself
- ✓ Asks before sending, paying, signing or deleting, enforced outside the AI

WHERE IT FALLS SHORT

- ✗ Exists only on personal Pro and Max plans so far
- ✗ Doesn't connect to Microsoft 365 yet
- ✗ Checks every few hours, so it can miss urgent items

YOUR DATA

- **Training:** On unless you turn off "Help improve Claude"
- **Who can see it:** Safety staff only; cloud computer access undisclosed
- **Deleting it:** Disconnecting an app isn't enough; reset memory by hand
- **Business plan:** No contract on how your data is handled (DPA) yet

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Memory, skills and routines from your Claude account
- **What's enforced vs guidance:** Approval gates are hard; your rules are guidance
- **Evidence:** Anthropic warns it may misread threads or misjudge urgency

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
Yes, opt-out	Opt-out training, no DPA	Manual memory reset	2025 switch to training default	View, edit, reset	Safety gates only

Sources: Anthropic Duet safety article (Oct 23, 2026) · Claude Help Center (Oct 2026) · Anthropic privacy policy (Sep 10, 2026) · Axios on the FTC (Oct 26, 2026) · Researched Oct 28, 2026

AIWhy

Claude Duet

Claude Duet is Anthropic's always-on personal assistant, launched October 21, 2026. It runs Claude Opus 5.5 on its own cloud computer, checks your connected accounts at set intervals, sends a morning brief and drafts things like invoice reminders for your review. Because it checks on a schedule, it can miss something urgent in between.

For now it's a US waitlist preview on personal plans only, from $20 to $200 a month, with no Microsoft 365 connection yet. Those plans train on your data by default, and there's no data processing agreement (DPA). Its password design is the best documented of the group: the AI never sees your passwords. Dan Shipper called it "the first personal agent I'd call a pro tool."

Switch off "Model Improvement" on day one, and keep client, HR and regulated data out.

SOURCES
Get started with Claude Duet, Claude Help Center: https://support.claude.com/en/articles/17631495-get-started-with-claude-duet
Claude pricing, Anthropic: https://claude.com/pricing
Claude Duet Safety, Privacy, and Security, Anthropic: https://claude.com/resources/articles/claude-duet-safety-privacy-and-security
Anthropic Privacy Policy: https://www.anthropic.com/legal/privacy
Is my data used for model training?, Anthropic Privacy Center: https://privacy.claude.com/en/articles/10023580-is-my-data-used-for-model-training
Vibe Check: Duet, Anthropic's Always-on Agent, Dan Shipper, Every: https://every.to/vibe-check/vibe-check-duet-anthropic-s-always-on-agent

OpenAI

OpenAI Dots

OpenAI · Launched Sep 29, 2026, rolling out · Pro $100–$500/mo; Business $100–125/seat

BEST FOR Delegate work in ChatGPT, Slack, Teams

FOR BUSINESS DATA

CAUTION

Use Business Premium at minimum; keep regulated and client-confidential data out.

WHAT IT DOES WELL

- ✓ Works from ChatGPT, Slack, Teams and voice calls
- ✓ Catches missed admin, like a misfiled email, without being asked
- ✓ Runs its own cloud computer you can take over anytime
- ✓ Hands passwords and money transfers back to you, every time

WHERE IT FALLS SHORT

- ✗ Hides its memories; deleting means wiping the whole dot
- ✗ Stalls on bot checks, CAPTCHAs and two-step login codes
- ✗ Strays from scope in 19.7% of ten-task chains, per OpenAI

YOUR DATA

- **Training:** On by default for Pro; off on Business plans
- **Who can see it:** Human safety reviewers, even after you opt out
- **Deleting it:** Delete the whole dot; copies stay in files and chats
- **Business plan:** Yes, a contract on how your data is handled (DPA)

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Chat, feedback and Custom Rules for each action
- **What's enforced vs guidance:** A second AI checks actions; not a guarantee
- **Evidence:** OpenAI says it "can make mistakes, including when following your rules"

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
Pro: yes	Business: no training, DPA	Delete whole dot	2023 chat and payment leak	Can't view memories	AI-judged gates

Sources: OpenAI dots safety post (Sep 29, 2026) · OpenAI dots privacy FAQ (Oct 2026) · GPT-6 Astra system card (Sep 29, 2026) · Scientific American (Oct 9, 2026) · Researched Oct 28, 2026

AIWhy

OpenAI Dots

Dots is OpenAI's always-on assistant inside ChatGPT, launched September 29, 2026. Each person gets one dot, working on its own cloud computer. You can reach it in ChatGPT, Slack and Teams or by voice call, and it can work with files on your computer if you turn that on.

Pro plans ($100 to $500 a month) train on your data by default. Business Premium seats ($100 to $125) don't, and they're HIPAA-eligible, so that's the lowest tier we'd use. Eligible Enterprise workspaces get dots switched on by default from November 18.

Memory is the weak spot. You can't view or delete a single memory; clearing one means deleting the whole dot. Users report a dot that couldn't reopen a spreadsheet it had saved, though nobody has verified that.

SOURCES

Introducing dots, OpenAI: https://openai.com/index/introducing-dots/
Getting started with your dot, OpenAI Help Center: https://help.openai.com/en/articles/20001530-getting-started-with-your-dot
Dots privacy, security, and safety FAQs, OpenAI Help Center: https://help.openai.com/en/articles/20001529-dots-privacy-security-and-safety-faqs
Managing dots in ChatGPT workspaces, OpenAI Help Center: https://help.openai.com/en/articles/20001554-managing-dots-in-chatgpt-workspaces
ChatGPT Business overview, OpenAI Help Center: https://help.openai.com/en/articles/8792828-chatgpt-business-overview

Meta Muse

Meta · Live in the US since Sep 8, 2026 · Free; Power $20/mo; Maximum $100/mo

BEST FOR Run personal life across Meta apps

FOR BUSINESS DATA

AVOID

On every plan, free to $100; revisit once Meta's promised lockout ships and is audited.

WHAT IT DOES WELL

- ✓ Works inside WhatsApp, Instagram, Facebook and Marketplace
- ✓ Handles Marketplace buying and seller messages where rivals get blocked
- ✓ Checks every app action through a separate gatekeeper, Sentinel
- ✓ Starts free and needs no technical skills to use

WHERE IT FALLS SHORT

- ✗ Gets blocked by Amazon; its model isn't top tier
- ✗ Has no business plan, team seats or admin console
- ✗ Keeps hourly notes in your account on everyone you mention

YOUR DATA

- **Training:** On by default on every tier; chats kept out of ads
- **Who can see it:** Meta can enter your cloud computer when "needed"
- **Deleting it:** Meta admits Muse "may still remember" what you delete
- **Business plan:** None; no contract on how your data is handled (DPA)

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Editable memory files, plus skills saved from chat
- **What's enforced vs guidance:** Sentinel's permission gates are hard; memory is guidance
- **Evidence:** "Allow Always" let Muse share a seller's home address

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
Yes, all tiers	Trains; Meta can access VM	"May still remember"	$5B FTC, €1.2B GDPR	Editable files	Sentinel gates only

Sources: Meta Help Center on Muse (Oct 2026) · Meta, "How we built safety into Muse" (Sep 8, 2026) · TIME (Oct 6, 2026) · Lenny's Newsletter (Sep 16, 2026) · Researched Oct 28, 2026

AIWhy

Meta Muse

Muse is Meta's personal AI agent, launched September 8, 2026, inside WhatsApp, Instagram, Facebook and Marketplace. It's live in the US, Canada, Japan and Singapore, free or at $20 and $100 a month.

The data terms are why we say avoid. Muse trains on your conversations by default on every plan, paid ones included. Meta says content you delete "may still" be remembered, Meta can get into the cloud computer Muse works on, and Muse keeps hourly notes on the people you mention. In October about 6,800 shared Muse chats were exposed. There's no business plan, no data processing agreement (DPA) and no SOC 2 report.

We'd look again once Meta's promised lockout of its own access ships and has been audited.

SOURCES
Introducing Muse: The World's First Personal AI Agent Built for Everyone, Meta Newsroom: https://about.fb.com/news/2026/09/introducing-muse-personal-ai-agent/
How to manage your Muse data, Meta Help Center: https://www.meta.com/help/artificial-intelligence/2225571704857152/
Security and safety for AI agents: our approach with Muse, Meta: https://muse.ai/blog/security-and-safety-for-ai-agents-our-approach-with-muse
Meta debuts Muse Wonder AI model in latest test of users' willingness to pay, CNBC: https://www.cnbc.com/2026/10/12/meta-muse-wonder-ai-model-agent.html
Meta's Muse AI Agent Is Building a Dossier On You, TIME: https://time.com/article/2026/10/06/meta-muse-ai-agent-privacy/
Meta halts Muse share links after 6,800 turn up in public DNS records, PPC Land: https://ppc.land/meta-halts-muse-share-links-after-6-800-turn-up-in-public-dns-records/

Instinct

Instinct

Spear Street Technology · Invite-only beta since Mar 2026 · Free during beta

BEST FOR **Outsource errands by text**

FOR BUSINESS DATA

AVOID

No business plan exists; keep company email, client data and payment cards off the free beta.

WHAT IT DOES WELL

- ✓ Takes requests by text or phone call, with no app.
- ✓ Books travel and restaurants, orders groceries, cancels subscriptions.
- ✓ Follows up on open tasks without being asked twice.
- ✓ Pays with one-time cards for amounts you approve.

WHERE IT FALLS SHORT

- ✗ Sent an email for a tester without asking first.
- ✗ Binds you to deals it makes, but disclaims its mistakes.
- ✗ Offers no business plan, admin console or audit log.

YOUR DATA

- **Training:** On by default; opting out only covers future data.
- **Who can see it:** Undisclosed; its bot said emails were plain text.
- **Deleting it:** Disconnecting doesn't delete; request deletion separately.
- **Business plan:** None, and no contract on how data is handled (DPA).

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** In one long chat; it stores your preferences.
- **What's enforced vs guidance:** No settings you control; approvals are promises.
- **Evidence:** No memory viewer; it acted once without approval.

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
● Yes	● Kept using data after disconnect	● Disconnect ≠ delete	● License grab, refused deletion	● No viewer	● Acted unapproved

Sources: TechCrunch (Aug 24, 2026) · Instinct Terms and Privacy Policy (Aug 26, 2026) · Business Insider (Sep 24, 2026) · Fortune (Sep 30, 2026) · Researched Oct 28, 2026

AIWhy

Instinct

Instinct, from Spear Street Technology, is an assistant you text or call, with no app to learn. It books travel and restaurants, orders groceries, cancels subscriptions and can call businesses for you. The terms say personal use only, and there's no business tier. It has been free in beta, and fees may apply later.

Training is on by default, and opting out only covers data from then on. The company publishes no retention periods, doesn't name its AI model providers and has no SOC 2 report.

Watch the money. Payments use one-time Link cards you approve, but a new Link balance (October 26) lets purchases under a limit you set skip approval. Keep company email, client data and payment cards off it.

SOURCES

Instinct homepage, Spear Street Technology: https://www.instinct.com
Instinct Terms of Service: https://www.instinct.com/terms
Instinct Privacy Policy: https://www.instinct.com/privacy
Set up a balance for agent wallet purchases, Link Support: https://support.link.com/questions/set-up-a-balance-for-agent-wallet-purchases
Instinct brings its group AI agent to iMessage for all users, RuntimeWire: https://runtimewire.com/article/instinct-group-ai-agent-imessage
Instinct Isn't Working On WhatsApp For Many Users, Reportedly Over A Billing Issue, OfficeChai: https://officechai.com/ai/instinct-isnt-working-on-whatsapp-for-many-users-reportedly-over-a-billing-issue/

Grok Bot

SpaceXAI / Cursor · Beta since Aug 2026 · $30–$200/mo plus usage overage

BEST FOR **Staff a roster of specialist bots**

FOR BUSINESS DATA

CAUTION

Only on Cursor Enterprise with Privacy Mode enforced; keep customer and financial data off.

WHAT IT DOES WELL

- ✓ Coordinates named bots that message each other and share work
- ✓ Drives any website as you, even ones without connectors
- ✓ Drafts a skill from a 10-minute demo you record
- ✓ Keeps passwords and saved cards hidden from the AI

WHERE IT FALLS SHORT

- ✗ Shares one computer and every login across all your bots
- ✗ Picks its own AI model; your team's vendor allowlist doesn't apply
- ✗ Shows only the last 20 runs, with no step-by-step failure view

YOUR DATA

- **Training:** Off only with Cursor Privacy Mode; individual default undocumented
- **Who can see it:** Staff access undisclosed; flagged data may be stored
- **Deleting it:** Delete your whole Cursor account; bot deletion leaves files
- **Business plan:** Cursor Enterprise; contract on how your data is handled (DPA)

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Chat rules, skills, recorded demos and scheduled routines
- **What's enforced vs guidance:** Approval rules are hard; skills and memory aren't
- **Evidence:** No memory screen; each bot decides what to save

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
Privacy Mode off	Privacy Mode; affiliate sharing	No retention policy	370K chats publicly indexed	Chat only	Approvals only

Sources: xAI Grok Bot docs (Oct 2026) · Cursor data-use page (Sep 3, 2026) · Irish DPC (Apr 2025, Feb 2026) · Forbes on leaked Grok chats (Aug 2025) · Researched Oct 28, 2026

AIWhy

Grok Bot

Grok Bot comes from SpaceXAI and runs on Cursor, which SpaceX bought in August 2026, so you sign in with a Cursor account. Instead of one assistant you get a roster of specialist bots, and version 1.0 (November 2) added shared Team Bots in Slack and Teams and bots that join calls.

Plans run $30 to $200 a month plus overage, and the spending cap doesn't stop a task already running. Over Halloween weekend Musk tightened usage limits, calling the costs "unsustainably high."

Two risks drive our caution. All of one user's bots share a single cloud computer and every login on it. And unless you add "Ask first" rules, a saved card fills in at checkout with no fresh approval. Stick to team plans with Cursor's Privacy Mode on.

SOURCES
Introducing Grok Bot 1.0, SpaceXAI: https://x.ai/news/grok-bot-1-0
Grok Bot FAQ, SpaceXSI docs: https://docs.x.ai/grok-bot/faq
Grok Bot: Approvals, security and privacy, SpaceXSI docs: https://docs.x.ai/grok-bot/approvals-security-and-privacy
The X Subscription pricing, SpaceXSI docs: https://docs.x.ai/x-subscription/pricing
Data Use and Privacy Overview, Cursor: https://cursor.com/data-use
Elon Musk throttles Grok Bot limits due to "unsustainably high" costs, TechCrunch: https://techcrunch.com/2026/11/02/elon-musk-throttles-grok-bot-limits-due-to-unsustainably-high-costs/

n8n Nod

n8n Labs · Research preview since Nov 10, 2026 · Free; you pay your AI provider

BEST FOR A personal assistant on your own computer

FOR BUSINESS DATA

CAUTION

Tasks go to your AI provider and app connections run through n8n's hosted service; a preview with no security terms yet.

WHAT IT DOES WELL

- ✓ Asks before it sends, orders, pays, books or deletes, enforced in code
- ✓ Shows each task's cost, with a monthly spending cap
- ✓ Memory you can view, edit, end-date or forget
- ✓ Web pages and emails can't give it orders

WHERE IT FALLS SHORT

- ✗ Stops when your laptop sleeps
- ✗ Slow on cheap models: one tester waited 14 minutes for a first reply
- ✗ Whether your words cover an action is judged by an AI that "can be wrong"

YOUR DATA

- **Training:** Depends on the AI provider you choose; n8n says nothing
- **Who can see it:** Your AI provider sees each task; usage analytics on by default
- **Deleting it:** TBD; analytics already sent stay sent
- **Business plan:** None yet: no DPA or SOC 2 in the preview

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Do a task once; it saves a skill you then check and edit
- **What's enforced vs guidance:** Approval rules are code; skills and memory are guidance
- **Evidence:** A memory page shows each note's date and source

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
● Depends on your AI provider	● Stays on your machine; analytics on	● TBD	● TBD: too new	● View, edit, forget	● Approvals in code

Sources: nod.n8n.io and n8n blog (Nov 10, 2026) · Nod docs and README · MindStudio (Nov 12, 2026) · DEV Community (Nov 11, 2026) · Researched Nov 12, 2026; laid out Nov 17, 2026

AIWhy

n8n Nod

Nod is a free, open-source personal assistant released November 10, 2026 by n8n Labs, an experiments team inside n8n. Despite the name, it isn't built on n8n workflows. It runs on your own computer or server, and you pay only your AI provider. It stops when your laptop sleeps, so always-on means a Mac mini or a server.

Running it yourself doesn't keep your data at home. Your AI provider sees each task, app connections run through a service n8n hosts, and usage analytics are on by default. It asks before it sends, pays or deletes, but the check on whether your words covered an action is itself an AI, which n8n says "can be wrong." As a research preview it has no DPA or SOC 2 report yet.

SOURCES

Meet Nod, n8n blog: https://blog.n8n.io/meet-nod
Nod homepage, n8n: https://nod.n8n.io/
n8n-io/nod README, GitHub: https://github.com/n8n-io/nod
When Nod asks before acting, Nod docs: https://nod.n8n.io/docs/trust/approvals/
How memory works, Nod docs: https://nod.n8n.io/docs/use/memory/
Analytics, Nod docs: https://nod.n8n.io/docs/deploy/analytics/

Viktor

Zeta AI · Public since Feb 2026 · Free to start; team plans from $50/mo

BEST FOR Add an AI teammate in Slack

FOR BUSINESS DATA

CAUTION

Non-sensitive work only on Team; Enterprise nears safe once its full audit lands.

WHAT IT DOES WELL

- ✓ Joins Slack or Teams like a coworker you @mention.
- ✓ Writes and runs code to build reports, dashboards and apps.
- ✓ Runs scheduled jobs and proposes work you didn't ask for.
- ✓ Bills the whole team from pooled credits, with no seat fee.

WHERE IT FALLS SHORT

- ✗ Shares one agent across the whole team; no private mode yet.
- ✗ Burns credits unpredictably, even on failed runs.
- ✗ Holds only a first-stage security audit (SOC 2 Type 1).

YOUR DATA

- **Training:** Never, on any plan; its AI providers are barred too.
- **Who can see it:** The team's shared agent; staff access is undisclosed.
- **Deleting it:** Self-serve; gone in about 30 days, backups 35.
- **Business plan:** Every plan gets a public data contract (DPA).

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Correct it in chat; it saves a skill.
- **What's enforced vs guidance:** Approval gates enforced; "Always/Never" rules are only guidance.
- **Evidence:** Its help pages cover "Viktor forgot my preference."

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
● Never	● Never trains, public DPA	● 30-day purge	● None found (young company)	● Via chat only	● Gates; skills drift

Sources: Viktor Privacy Policy and DPA (Oct 2026) · Viktor Security page (Oct 2026) · Fortune (May 19, 2026) · Viktor Academy, Skills & Memory (Oct 2026) · Researched Oct 28, 2026

AIWhy

Viktor

Viktor, from Zeta AI, is one shared AI coworker for a whole Slack or Microsoft Teams workspace. A private mode for each person is still on the roadmap. You can pick the AI model, and Claude Fable 5.2 became the top preset on November 10.

Team plans start at $50 a month for credits the whole team shares. Credits burn fast, and failed runs cost credits too.

Its data terms are the best of the startups: no training on any plan, a published data processing agreement (DPA), and documented deletion. Two things keep it at caution. One click on "always approve this tool" removes that tool's approval gate, and so far it has only a first-stage security audit (SOC 2 Type 1). Enterprise gets close to safe once the full audit is done.

SOURCES

Viktor, homepage and FAQ: https://viktor.com
Viktor, Pricing: https://viktor.com/pricing
Viktor, Security: https://www.viktor.com/security
Viktor (Zeta AI, Inc.), Privacy Policy: https://viktor.com/legal/privacy
Viktor, Changelog: https://viktor.com/changelog
Trustpilot, Viktor reviews: https://www.trustpilot.com/review/viktor.com?stars=1&stars=2

perplexity

Perplexity Teammates

Perplexity · Research preview since Nov 5, 2026 · Needs Max, $200/mo; runs on credits

BEST FOR Role agents working in Slack, GitHub and Linear

FOR BUSINESS DATA

AVOID

Max trains on your data by default and past data can't be deleted; wait until Teammates reaches Enterprise.

WHAT IT DOES WELL

- ✓ Runs in the background and pages you only for decisions
- ✓ Works in Slack, GitHub and Linear, where the team already is
- ✓ Each teammate gets its own computer; logins hidden from it
- ✓ You can take over its desktop mid-task

WHERE IT FALLS SHORT

- ✗ Acts with one owner's permissions, even in a shared channel
- ✗ No approval rules you can set; it asks at its own checkpoints
- ✗ Credits burn fast: a daily research task can use the month's allotment
- ✗ Every success story so far is a Perplexity demo

YOUR DATA

- **Training:** On by default on Max; opting out works going forward only
- **Who can see it:** TBD: staff access and human review undisclosed
- **Deleting it:** Past training data "cannot be deleted"
- **Business plan:** Enterprise never trains, but Teammates isn't on by default

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** A written job description per teammate, plus your feedback
- **What's enforced vs guidance:** Guidance: "every time without asking" is soft memory
- **Evidence:** TBD: no teammate memory screen documented

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
On for Max	Enterprise only	Past data stays	Class action; Comet prompt injection	TBD	Its own checkpoints

Sources: Perplexity Teammates help center and product page · launch blog (Nov 5, 2026) · Perplexity privacy help · CellCog (Nov 5, 2026) · Researched Nov 12, 2026; laid out Nov 17, 2026

AIWhy

Perplexity Teammates

Teammates, announced about November 5, 2026, are Perplexity's agents for ongoing work, one per role, in Slack, GitHub and Linear. Each one belongs to a single Perplexity Max subscriber ($200 a month) and acts with that person's permissions, even in a shared channel. It's a research preview.

Credits burn fast. A daily research task can use up a month's allotment, and tasks pause when credits run out.

Max trains on your data by default, and Perplexity says past data "cannot be deleted." Enterprise never trains, but it doesn't get Teammates by default. You can't set approval rules either; a teammate asks at its own checkpoints. Every success story so far is a Perplexity demo. We'd wait until Teammates reaches Enterprise.

SOURCES

Get started with Perplexity teammates, Perplexity Help Center: https://www.perplexity.ai/help-center/en/articles/2026102601-perplexity-teammates
Teammates product page, Perplexity: https://www.perplexity.ai/hub/products/teammates
What is Computer?, Perplexity Help Center: https://www.perplexity.ai/help-center/en/articles/13837784-what-is-computer
Computer for Enterprise, Perplexity Help Center: https://www.perplexity.ai/help-center/en/articles/13901210-computer-for-enterprise
Data collection at Perplexity, Perplexity Help Center: https://www.perplexity.ai/help-center/en/articles/11564572-data-collection-at-perplexity
Perplexity Computer Can Now Run Fully Locally on Mac, MacRumors: https://www.macrumors.com/2026/10/20/perplexity-computer-local-mac/

Microsoft Copilot Autopilot

Microsoft · Private preview since Sep 2026 · $21–30/seat plus usage credits

BEST FOR **Automate Microsoft 365 work**

FOR BUSINESS DATA

SAFE

On Microsoft 365 Copilot work accounts, once it leaves private preview and SharePoint is cleaned up.

WHAT IT DOES WELL

- ✓ Works in Teams and Outlook like a coworker you @mention.
- ✓ Acts under its own work identity, mailbox and calendar.
- ✓ Watches channels, follows up threads, resumes projects days later.
- ✓ Checks your company's data-loss rules before it sends anything.

WHERE IT FALLS SHORT

- ✗ Can't be bought yet; private preview, no launch date.
- ✗ Charges metered usage credits at unpublished rates.
- ✗ Acts on whatever your SharePoint permissions already overshare.

YOUR DATA

- **Training:** Never on work accounts; no human review either.
- **Who can see it:** Undisclosed whether staff can reach the agent's computer.
- **Deleting it:** You can delete history; agent memory deletion undisclosed.
- **Business plan:** Data contract (DPA), health-data (HIPAA) and EU terms.

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Give it a name, role, goal and routines.
- **What's enforced vs guidance:** Access rules enforced; its methods are guidance.
- **Evidence:** None yet; no independent hands-on tests published.

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
● Never (work accounts)	● No training, DPA, no review	● Undisclosed	● None found in research	● Undocumented	● IT policies enforced

Sources: Microsoft Blog, Autopilot launch (Sep 25, 2026) · Microsoft Learn, Copilot privacy (Oct 26, 2026) · Message Center MC1479277 (Oct 12, 2026) · Computerworld on Forrester (Jun 2026) · Researched Oct 28, 2026

AIWhy

Microsoft Copilot Autopilot

Autopilot is Microsoft's always-on agent in the new Copilot app, announced September 25, 2026. It works in Teams and Outlook like a coworker you @mention, with its own work identity, mailbox and calendar, and it can pick projects back up days later.

It's rolling out in November to Frontier, Microsoft's early-access program, with no general release date. You need a Microsoft 365 Copilot license; our October research put that at $21 to $30 a seat plus usage credits. Microsoft's Ignite conference runs November 17 to 20.

On work accounts the data terms are the strongest in the field. The work is on your side. Autopilot acts on whatever your people can already see, so clean up your SharePoint permissions first. Anything SharePoint overshares, it can use too.

SOURCES
What's New in Microsoft Copilot, October 2026, Microsoft Tech Community: https://techcommunity.microsoft.com/blog/microsoft-copilot-blog/what%E2%80%99s-new-in-microsoft-copilot--october-2026/4568528
Microsoft Ignite 2026 Dates: 17–20 Nov, Solv Systems: https://solv-systems.com/resources/microsoft-ignite-2026-preview
Copilot November 2026 Update, FreakyJolly: https://freakyjolly.com/microsoft-copilot-november-2026-update/

OpenClaw

OpenClaw Foundation, open source · OpenClaw 2.0 since Aug 2026 · Free software + your AI and hosting bills

BEST FOR Self-host one agent across every chat

FOR BUSINESS DATA

AVOID

Self-hosted, without IT staff to harden it; caution only if a technical operator locks it down.

WHAT IT DOES WELL

- ✓ Reaches you in about 30 chat apps, from WhatsApp to Slack
- ✓ Runs on your own machine with any AI model you choose
- ✓ Keeps its memory in plain files you can read and edit
- ✓ Runs scheduled jobs and works your computer's files and browser

WHERE IT FALLS SHORT

- ✗ Ships with safety off: no locked-down workspace (sandbox), no approvals
- ✗ Let 341 malicious add-ons (skills) into its public skill store
- ✗ Needs real IT skill to lock down; no vendor support

YOUR DATA

- **Training:** Project collects nothing; pay-per-use AI accounts (API keys) don't train
- **Who can see it:** You, your hosting company, and any add-on you install
- **Deleting it:** Delete the files and it's gone; AI provider may keep 30 days
- **Business plan:** None: no security audit (SOC 2), data-handling contract (DPA) or support

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Plain text files it rereads each session, plus self-written skills
- **What's enforced vs guidance:** Memory is guidance; approvals and sandbox are firm but off by default
- **Evidence:** Its own docs say it plainly: "Memory does not enforce policy"

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
● No (API keys)	● No vendor; your files	● Your files	● None (security excluded)	● Plain files	● Off by default

Sources: OpenClaw docs (Oct 2026) · The Hacker News, ClawHub malware (Feb 2026) · SecurityScorecard (Feb 2026) · Trail of Bits audit (Sep 2026) · Researched Oct 28, 2026

AIWhy

OpenClaw

OpenClaw is free, open-source software that runs an always-on agent on your own machine, with full access to your files. It reaches you in about 30 chat apps and ships a release every few days. Multiplayer, launched November 5, lets several people share one agent, and the launch post describes no permission controls for it.

You pay for the AI model and the hosting. There's no vendor support, no security audit (SOC 2) and no data processing agreement (DPA).

Its security record is the worst in the field. Malicious add-ons (skills) in its public store passed 575 by November 10, up from 341 in February. It ships with no locked-down workspace (sandbox) and no approvals until you turn them on. Without IT staff to lock it down, we'd stay away.

SOURCES
OpenClaw GitHub README: https://github.com/openclaw/openclaw
OpenClaw GitHub Releases: https://github.com/openclaw/openclaw/releases
OpenClaw blog, "Shattering Silos with OpenClaw Multiplayer": https://openclaw.ai/blog/shattering-silos-with-openclaw-multiplayer
OpenClaw blog index: https://openclaw.ai/blog
OpenClaw docs, memory: https://docs.openclaw.ai/concepts/memory
Help Net Security on Acronis TRU Cyber Pulse 2026: https://www.helpnetsecurity.com/2026/11/10/acronis-ai-agent-cyberattacks-report/

Hermes Agent

Nous Research, open source · Released Feb 2026 · Free software + ~$2–60/mo in AI use; hosted $20–$200/mo

BEST FOR Self-host an agent that learns routines

FOR BUSINESS DATA

CAUTION

Self-hosted with a business API key and skill approvals on; avoid Nous's hosted plans for client data.

WHAT IT DOES WELL

- ✓ Writes its own how-to notes (skills) after solving a problem
- ✓ Searches every past conversation to recall how you worked
- ✓ Reaches you in about 30 chat apps and runs scheduled jobs
- ✓ Runs on your hardware with any AI model, including local ones

WHERE IT FALLS SHORT

- ✗ Rewrites and deletes its own skills without asking, by default
- ✗ Lets a helper AI auto-approve "low-risk" commands by default
- ✗ Has ~10 published security flaws (CVEs) in 2026; an audit sits unanswered

YOUR DATA

- **Training:** Self-hosted: no. Nous's hosted plans: yes, unless Privacy Mode is on
- **Who can see it:** You and your host; Nous Cloud staff access isn't disclosed
- **Deleting it:** Delete the files; Nous deletes hosted accounts within six months
- **Business plan:** No security audit (SOC 2) or data-handling contract (DPA), any tier

DOES IT REMEMBER WHAT YOU TEACH?

- **How you teach it:** Correct it once; it saves the fix as a skill file
- **What's enforced vs guidance:** Skills and memory are guidance; command blocklists and allowlists are firm
- **Evidence:** An audit rated self-written skills Critical: planted instructions can stick

TRAINS BY DEFAULT	DATA PROTECTIONS	DELETE DELETES	DATA-HANDLING RECORD	MEMORY VISIBLE	RULES ENFORCED
● No (self-hosted)	● No vendor; your files	● Your files	● None found	● Plain files	● Rewrites own skills

Sources: Hermes Agent docs (Oct 2026) · Nous Research terms (Sep 2026) · GitHub security audit #7826 (Apr 2026) · BleepingComputer (Jul 2026) · Researched Oct 28, 2026

AIWhy

Hermes Agent

Hermes Agent is open-source software from Nous Research, first released in February 2026. It runs on your own hardware with any AI model and reaches you in about 30 chat apps. After it solves a problem, it writes its own skills, which are how-to notes it reuses.

Expect about $2 to $60 a month in AI use, or $20 to $200 on Nous's hosted plans. Self-hosted, nothing trains on your data; the hosted plans do unless you turn on Privacy Mode.

Those self-written skills are the risk. By default it rewrites and deletes them without asking, and an audit rated them critical because a planted instruction can persist. It also has about 10 published security flaws (CVEs) in 2026. Turn on skill write approval, and keep client data off the hosted plans.

SOURCES

Hermes Agent GitHub README: https://github.com/NousResearch/hermes-agent
Hermes Agent releases, GitHub: https://github.com/NousResearch/hermes-agent/releases
Nous Research, hermes-agent SECURITY.md: https://github.com/NousResearch/hermes-agent/blob/main/SECURITY.md
Hermes Agent docs, memory: https://hermes-agent.nousresearch.com/docs/user-guide/features/memory
Nous Research opens the official Hermes Agent mobile beta, AICoder: https://aicoder.com/news/news-20261031-nous-hermes-agent-mobile-beta-ios-testflight-android-firebase-own-gateway
Nous Research Ships Hermes Agent v0.23.0 The Hearth, AICoder: https://www.aicoder.com/news/news-20261105-nous-hermes-agent-v0-23-0-hearth-webapp-unified-gateway

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

AIWhy