

BRIEF IT / CHECK IT / KEEP IT SAFE

The 2026 AI Starter

AI assistants can now take a goal and hand back finished work.
This is how to hand that work over and stay in charge of it.

01

Hand it the whole job, in a six-part brief.

02

Run five checks before you pass anything on.

03

Set how far it may go, in the tool itself.

ALSO INSIDE

CLAUDE, CHATGPT OR GEMINI

CONNECTORS AND MCP

YOUR FIRST WEEK

12 STARTER BRIEFS

A PIN-UP SHEET

For anyone at work who uses Claude, ChatGPT or Gemini.

Neosaipien is not affiliated with Anthropic, OpenAI or Google.

Read it once, then try one habit a day

Claude, ChatGPT and Gemini can now take a goal and do the work, depending on your plan: researching, drafting, editing files, sending email. What matters now is less the wording and more how you hand the work over, check it and set its limits. This guide is the short version, with my own pick of which assistant to use for what.

01

Read it once

Pages 3 to 11: what changed, which assistant for what, where to click in each, and the three habits.

02

One habit a day

Page 12 suggests a first week: one small job a day, start to end, using work you'd do anyway.

03

Keep the tools

Twelve starter briefs and a sheet to pin up. The briefs also come as a text file, so you can copy each one whole.

CONTENTS

What changed this year	03
Claude, ChatGPT or Gemini	04
Where to click, start to end	05
Which one is this job?	06
Habit 1: Brief it	07
A worked example	08
Habit 2: Check it	09
Habit 3: Keep it safe	10
Connectors and MCP	11
Your first week	12
Twelve starter briefs	13
Pin this up	15
Checked against	16
Where to go next	18

BEFORE YOU PASTE OR CONNECT

Check your workplace's rules before you paste anything into an AI tool, or connect it to your email or files. Customer details, staff records, contracts and figures may not be allowed. This guide is general guidance, not legal advice, and your organisation's own rules come first.

WHICHEVER ASSISTANT YOU USE

It works with Claude, ChatGPT or Gemini: use the one your workplace approves. Page 4 compares the three, with my recommendation. What each can do depends on your plan and changes every month, so every product fact here is dated as of October 2026 and was read on the maker's own pages.

BY THE END OF YOUR FIRST WEEK, YOU CAN

1

Pick the assistant, the model and the kind of tool a job needs.

2

Hand a job over in a six-part brief that a colleague could follow.

3

Check what comes back with five checks, outside the chat.

4

Connect an app safely, and set what it may do without you.

From answers to finished work

Four changes, as of October 2026, each read on its source's own page.

01

It takes a goal and does the work

On 9 Jul 2026 OpenAI launched ChatGPT Work, an agent for longer tasks and finished files, in a desktop app that “brings Chat, Work, and Codex together”. On 16 Sep 2026 Anthropic merged Cowork into ordinary Claude chat: “Users no longer have to choose which mode a task belongs in.” Gemini’s agent is tasks in Gemini, for personal accounts on Google AI Pro or Ultra for now. OpenAI’s advice: “Describe the outcome you want.”

OpenAI, ChatGPT release notes, 9 Jul 2026, and ChatGPT Work and Codex; Anthropic, 16 Sep 2026; Google, tasks in Gemini Apps

02

New models, picked by the job

Anthropic released Claude Fable 5.1 on 1 Sep 2026 and Claude Opus 5.5 on 22 Sep 2026; OpenAI released GPT-6 Astra on 3 Sep 2026; Google released Gemini 4 Argon on 13 Oct 2026. These are the makers’ models for hard, long work. A heavier model or a higher thinking level is slower and uses more of your limit, so start with the default and move up only when it matters (page 5).

Anthropic, 1 and 22 Sep 2026; OpenAI, 3 Sep 2026; Google, 13 Oct 2026

03

Wording isn’t a safety plan

The UK’s National Cyber Security Centre, writing about agents on 20 Aug 2026, says to tell an agent “both what it should do and what it should not do”, and also: “do not rely on prompting alone.” What an agent may touch, send or run on a schedule is set in the tool, not only in your words. Anthropic’s engineers agree that “the exact formatting of prompts is likely becoming less important as models become more capable.”

UK NCSC, Managing the cyber risk of agentic AI, 20 Aug 2026; Anthropic engineering blog, 29 Sep 2025

04

Skills are replacing the prompt pack

A skill is a saved how-to: a file called SKILL.md with a name, a one-line description and the steps. The format was “originally developed by Anthropic, released as an open standard”, and Claude, ChatGPT and Gemini all accept it now. Custom GPTs are moving to plugins, and Gems become skills from Nov 2026 for personal accounts and in 2027 for work and school ones.

Agent Skills overview; OpenAI release notes, 11 Sep 2026; Google, the transition from Gems to skills

MAGIC WORDS WERE NEVER THE TRICK

Offer the AI a tip, or even threaten it? In a 2025 study of 2024 and 2025 models, researchers at Wharton tested threats and tips on nearly three hundred hard questions, running each one again and again. Overall, it made no real difference.

Wharton Generative AI Labs, Prompting Science Report 3, 1 Aug 2025

Prompt wording still helps, and it is now the smallest of the skills.

The bigger ones fill the rest of this guide: handing the work over well, checking what comes back, and keeping it inside safe limits.

Which assistant for what

All three now take a goal and hand back finished work, and each is strongest in different places. The middle columns are from the makers' own pages, as of October 2026. The last is my pick.

THE JOB	CLAUDE	CHATGPT	GEMINI	MY PICK
Writing and documents	Clear business writing (Anthropic's claim for Opus 5.5); finished Word, PDF or Google Docs files	Fast drafts in Chat; reports from Work; ChatGPT for Word, even on Free	Writing inside Gmail and Docs, at a company on Google Workspace	Claude
Spreadsheets	Claude for Excel: answers point to the cells, edits keep the formulas, and it warns before overwriting	Add-ins for Excel and Google Sheets; the Data plugin	Gemini in Sheets, on Google Workspace	Claude
Slides	Claude Slides on every plan; Claude for PowerPoint	Following your company's templates (OpenAI's claim for GPT-6 Astra)	Gemini in Slides	Claude, with ChatGPT close
Research	Mixes the web with the work tools you connect	The most control: edit its plan, and limit it to sites you name	Deep Research on the free plan; Gemini Notebook for your own files	ChatGPT; Gemini Notebook for your own files
Meetings	Give it the transcript or your notes	The Meetings plugin (beta, Mac) and audio uploads	"Take notes for me" in Google Meet	Gemini in Google Meet, otherwise ChatGPT
Email	The Gmail connector; Claude for Outlook	Gmail and Outlook Email apps; tasks that start on a new Gmail message	Gemini built into Gmail	Gemini in Gmail; Claude or ChatGPT for Outlook
Pictures and video	Designs, mockups and diagrams in Claude Design	ChatGPT Images 2.5	Nano Banana images; Gemini Omni video	ChatGPT or Gemini

Each maker publishes tests it wins, so this page rests on what each one offers you, not on its scores.

MY RECOMMENDATION

My recommendation: Claude first, ChatGPT second, Gemini third.

My own view, for most office work. It sits beside the honest table above, not instead of it.

1 CLAUDE

The simplest start to finish: one conversation for a question or a whole task, no mode to choose. It asks before acting by default. Docs, Slides and Design on every plan.

2 CHATGPT

A strong agent in Work: a cloud browser for sites you sign in to, and tasks that start when an email arrives. Plus voice, pictures and Sites.

3 GEMINI

The best fit if your company runs on Google Workspace, with a generous free plan. But its agent tasks and skills are for personal accounts for now.

WHERE ANOTHER ONE WINS, USE IT

Pictures and video: ChatGPT or Gemini. Google Meet notes and Gmail: Gemini. Routines that start on a new email or Slack message, and research where you control the sources: ChatGPT.

THE FIRST RULE

Use the assistant your company approves for company data. A preference, mine included, never beats your company's data rules.

Where to click, start to end

Each maker has its own words for the same idea. Here they are side by side, as of October 2026. A missing button usually means your plan, your country, a rollout on its way, or your admin.

THE JOB	CLAUDE	CHATGPT	GEMINI
Ask a quick question	Type in any conversation	Chat	Chat
Hand over a task	The same conversation: Cowork is built in	Choose Work, from the toggle or menu at the top	Click Task below the text box (personal account, Pro or Ultra)
Build software or a script	Claude Code, in the Code area	Codex, its own view in the app	Antigravity, in the Gemini Mac app
Keep a job's files together	A project: its chats, files and instructions	New project, in the sidebar	A notebook, in Gemini or Gemini Notebook
Tell it about you, once	Settings > Instructions for Claude	Settings > Personalization > Custom Instructions	Customize, in the sidebar
Save a how-to as a skill	Customize > Skills	Type @skill-creator and describe the task	Customize > Skills
Connect an app	Customize > Connectors	Settings > Plugins	Customize, or Settings > Connected Apps
Research with sources	Web search on, then "+" > Research	"+" > Deep research	Add Files > Deep Research
Repeat it on a schedule	Scheduled tasks	Scheduled Tasks, or start on a new email	A schedule in a task

CLAUDE CODE OR ORDINARY CLAUDE?

Claude Code is built for software. Cowork, built for other work, joined ordinary Claude on 16 Sep 2026, so the app has two places: Claude, and Code. Use Claude for documents, research, spreadsheets and repeat tasks; open Code only to build or change software, a website or a script. ChatGPT splits the same way: Chat for quick help, Work for a file or a done task, Codex for software.

WHICH MODEL WHEN

THE JOB	CLAUDE	CHATGPT	GEMINI
A quick question, a short rewrite	Haiku 5.5 or Sonnet 5.5	GPT-6 at Instant	Auto or Flash-Lite
Everyday writing, a first deck	Sonnet 5.5	GPT-6 at Medium	Flash
A careful report, a messy sheet	Opus 5.5	GPT-6 at High	Pro, thinking High
Hand over a task of many steps	Opus 5.5	Work on Default, or GPT-6.1 Sol	Task with Flash or Argon
The hardest, high-stakes work	Fable 5.1 (uses your limit fastest)	GPT-6 at Pro, or GPT-6 Astra in Work	Argon, thinking High

Where to pick: Claude, the model name by the send button. ChatGPT, GPT-6 in the model menu, then the Thinking slider. Gemini, the model under the chat box, then a thinking level.

Start with the default. Move up only when the answer misses something that matters; move down for speed, or to save your limit.

Not every job needs an agent

An agent is the newest tool, not always the right one. Start with the simplest that fits.

CHAT	DEEP RESEARCH	AGENT	AUTOMATION
One answer, draft or idea that you'll use yourself. You do any acting. FOR EXAMPLE Rewrite a paragraph for a customer. Explain a clause. OpenAI's own advice: "For quick lookups or short conversations, standard chat may be faster."	A question that needs many sources. It proposes a research plan you can edit, then hands back a report with citations (as of Oct 2026). FOR EXAMPLE Compare how five suppliers price the same service. Checking those sources becomes your job.	A goal whose steps you can't know in advance. It works across your files and apps and hands back finished work. FOR EXAMPLE Work out why sales dipped in one region last month. Give it a brief, a level and a check.	The same steps every time. If you can write the steps down in advance, you don't need an agent, you need a workflow. FOR EXAMPLE Send the same reminder every Friday. It's usually cheaper and faster, and it follows the same steps on Tuesday as it did on Monday.

The examples are illustrations.

HOW TO PICK, IN THIS ORDER

1	Is it one answer, draft or idea you'll use yourself?	CHAT
2	Does it need many sources, with citations?	DEEP RESEARCH
3	Can you write every step down in advance?	AUTOMATION
4	Otherwise, the next step depends on what it finds	AGENT

"using AI agents when decisions are needed, automation for routine workflows and assistants for simple retrieval"

Anushree Verma, Gartner, 25 Jun 2025

THREE QUESTIONS SET THE CEILING

Ask them of any job you hand to an agent, and take the lowest level any of them allows. The four levels are on page 10.

01 Can it be undone easily? If not: level 3 at most, and you approve each step in the tool.	02 Does it leave the team? A customer, a supplier, the public, money. If so: level 3 at most, and a person presses send.	03 Does it read outside content? Web pages, outside email, files others sent. If so: it must not also send, pay or delete without asking.
---	--	---

Brief it like a brilliant new employee

Anthropic's prompting guide says to picture a brilliant new employee, who doesn't yet know your standards or how you work. You wouldn't hand that person six words and walk away. You'd give them a brief, in six parts.

1 GOAL

What it is to achieve, why, and by when. One sentence.

2 DONE LOOKS LIKE

Who it's for, what it's for and what good looks like: the format, the length, the tone, and where it lands.

3 CONTEXT

The files, facts and examples to use, and which to trust. Anthropic's guide calls examples one of the most reliable ways to steer a model, and suggests three to five that differ.

4 MAY TOUCH, MUST NOT

What it may read, use and change, and what is off limits. Say both.

5 LIMITS

Time, steps, pages or spend: when to stop anyway.

6 STOP AND ASK


The moments it must come back to you: before anything it can't undo, when sources disagree, when it's unsure.

WHY YOUR BRIEF ONLY MAKES SENSE TO YOU

ONE STUDY, ABOUT SONGS


In a 1990 study, people tapped out well-known songs with a finger. On average, they expected the listener to name half. The listeners got one in forty.

The researchers' explanation: the tapper hears the whole tune in their head. The listener only hears taps.



THE TAPPERS' AVERAGE GUESS

1 IN 2



THE LISTENERS NAMED

1 IN 40

3 OF 120 SONGS

NEWTON, STANFORD, 1990

OUR COMPARISON

The study was about songs, not AI. The comparison is ours. When you brief an AI, you're the tapper. You hear the whole tune: who it's for, why it matters, what good looks like. The AI only gets what you typed.

So before you hand it over, put the tune into words: who it's for, and how they'll use it.

THE TEN SECOND TEST

Before you hand it over, imagine giving your brief to a colleague who knows nothing about the task. Could they follow it without coming back to ask you something? If they'd have questions, so does the AI. It may not stop to ask: it can fill each gap with a safe, typical guess instead. That's what Stop and ask is for.

We call it the ten second test because picturing that colleague takes about ten seconds. The name is ours. The idea is the golden rule in Anthropic's prompting guide.

From six words to a brief

BEFORE / 6 WORDS

Write an email about the delay.

A colleague would stop and ask you. The model usually just answers, so it fills every one of those gaps with a safe, typical choice. That's why the email sounds like it was written for anyone. Because it was.

LEFT OPEN, SO IT GUESSES

WHICH DELAY?

WHO IS IT GOING TO?

HOW BAD IS IT?

WHAT SHOULD THEY DO?

THE EMAIL IS AN EXAMPLE

AFTER / A LEVEL 2 BRIEF

Goal	Draft an email to [the clients whose order was due on Friday], telling them it will now arrive on [Tuesday], so they can plan and keep trusting us.
Done looks like	Who it's for: [clients who were promised delivery on Friday]. What it's for: to tell them the new date, own the mistake and keep their trust. What good looks like: under 120 words, plain language, the new date in the first line. One draft, saved in my Drafts folder.
Context	[The two-line note on what caused the delay], [the order list], and [last month's delay email] as an example of the tone.
May touch, must not	Read those three files and save one draft. Don't send anything, contact anyone or change the order records.
Limits	One draft, in about ten minutes.
Stop and ask	If the note and the order list give different dates, or if the email would need to promise a refund or a discount.

LEVEL 2 / DRAFT

WHY LEVEL 2

It drafts and never sends. Where your tool allows it, set that there too: drafting access only, or sending set to ask first. Then you read the draft and press send.

THE CHECK

The facts: the new date is one you can keep, and the cause it gives is true. The sign-off: delete any apology for something that didn't happen, then send it yourself.

The three things a good prompt always needed, who it's for, what it's for and what good looks like, now live inside Done looks like.

SAVE IT AS A SKILL

A brief you use every week can be saved as a skill, so your assistant follows it every time without retyping. As of Oct 2026, Claude, ChatGPT and Gemini all accept the same skill file, SKILL.md: in Claude and Gemini under Customize > Skills, and in ChatGPT by typing @skill-creator. Day 7 of your first week (page 12) does it.

Same request, now a brief. Every part you add takes a guess away.

Made up, to illustrate the method. It names no real company or person.

Finished-looking isn't the same as checked

A model writes one word at a time, by working out what's likely to come next, given everything on the page. And likely isn't the same as true. Quotes and links can be invented too, so the check happens outside the chat.

THE FIVE CHECKS

1	The job Did it do what the brief asked?	Hold it next to Done looks like: the right audience, format, length and deadline. Anything missing is a no, however polished the rest looks.
2	The facts Is every name, number, date and quote real?	Trace each one to its file or page, outside the chat. If you can't find it, it doesn't go in.
3	The steps How did it get there?	Read what it did, not just the summary at the end. Watch for files or sites you didn't mention.
4	The changes What is different now?	Look where it acted: drafts, sent items, edited files and their version history.
5	The sign-off Would you put your name to it?	You answer for what goes out, not the AI. If you wouldn't defend it in a meeting, it isn't done.

Asking it "are you sure?" isn't checking.
Check outside the chat.

When an AI tells you its work passed a check, ask what was checked.

50%

WORK TREND INDEX 2026

In Microsoft's survey of 20,000 workers using AI in 10 countries (published 5 May 2026), the two human skills rated most important as AI takes on more work were "quality control of AI output (50%)" and critical thinking (46%).

66%

UNIVERSITY OF MELBOURNE AND KPMG

In a study of more than 48,000 people in 47 countries (Nov 2024 to Jan 2025), many employees said they rely on AI output without evaluating its accuracy (66%) and have made mistakes in their work because of AI (56%).

SHORT OF TIME? CHECK THESE FIRST

- | | | | | | |
|---|---|---|-------------------------------------|---|------------------------------|
| 1 | Anything that leaves the team: a customer, a supplier, the public, money. | 2 | Every name, number, date and quote. | 3 | Anything it says it checked. |
|---|---|---|-------------------------------------|---|------------------------------|

BEFORE YOU FORWARD IT **Whoever briefed it or asked for it runs the checks. "The AI did it" is never the answer to "who checked this?"**

Set how far it may go, in the tool

An agent isn't a smarter AI. It's the same kind of AI, on a longer leash. And holding that leash is your job.

1 READ	2 DRAFT	3 ASK FIRST	4 ALONE
It looks and reports back. It changes nothing and sends nothing.	It makes drafts only you can see. Nothing goes out.	It stops before anything that sends, pays, deletes, books or changes a shared record, until you approve it.	It acts without asking, inside fixed limits, and you review the record afterwards.

These levels are our own rule of thumb, not a maker's setting: start low, go up a level after five clean runs, and down at once after a surprise. "Ask first" is set in the tool, not only typed in the brief. This starter stays at levels 1 and 2.

FOUR SETTINGS TO LOOK AT THIS WEEK As of Oct 2026, from the makers' own pages.

Connections What is it connected to, and can it write, or only read?	Claude: Customize > Connectors. ChatGPT: Settings > Plugins. Gemini: Customize, or Settings > Connected Apps. Each reaches only what your account can, and only what the connector allows (page 11).
Approvals What does it do without asking you first?	Claude asks before acting by default; each connector tool is Always allow, Needs approval or Blocked. ChatGPT's apps: Always ask, Allow read actions, Allow low-risk actions or Allow all actions. Gemini: an app's write actions always need your yes. While you learn: Needs approval, or Always ask.
Memory What does it remember, and can you clear it?	Claude: Settings > Memory, on by default for Free, Pro and Max; on Team and Enterprise, an owner turns it on and each person opts in. ChatGPT: Settings > Personalization > Memory. Gemini: Personal Intelligence stays off until you turn it on.
Schedules Does anything run while you're away?	Claude can start a report every Monday, and ChatGPT's Scheduled Tasks can also start on a new Gmail or Slack message. Anthropic: "Don't schedule tasks that access sensitive files, send messages on your behalf, make purchases, or take other actions that are difficult to undo."

Menus move and plans differ: if a setting isn't there, check your plan, or ask your admin.

IF IT DOES SOMETHING YOU DIDN'T EXPECT

OpenAI: "Stop the task if ChatGPT opens the wrong site or uses incorrect information."

Anthropic: "You remain responsible for all actions taken by Claude performed on your behalf."

In all three: Stop it, look where it acted, switch off the connector it used, and go down a level at once.

NEVER PASTE

OpenAI's own advice: "Never paste passwords, security codes, or payment details into the conversation." Google says the same for Gemini's tasks.

OUTSIDE WORDS, NO POWER

Anthropic says a prompt injection attack needs two things at once: the AI "can read information outside your trusted boundary, and can perform actions that could compromise the user." So an agent that reads outside content (the web, outside email, files others sent) may not also send, pay or delete without asking. Web pages, emails and documents can carry lines you can't see. **These steps cut the risk. They don't end it.**

Connectors and MCP, in plain words

A connector lets your assistant reach another app, such as your calendar, files or email, signed in as you, so it can look things up in your own work and act on it.

MCP

Model Context Protocol, a shared standard for plugging AI into other software: “Think of MCP like a USB-C port for AI applications.”

MCP SERVER

The other app’s doorway. It lists what an AI may do there: search files, read a message, create a record.

CONNECTOR

The ready-made cable to that doorway, signed in as you. ChatGPT says app; Gemini, connected app.

PLUGIN

A box that holds connectors and skills together, in Claude and ChatGPT.

SKILL OR CONNECTOR?

Anthropic: “MCP connectors give Claude access to a tool or data source, and skills teach it how to do a task.” A skill when instructions are enough; a connector when it needs live data or has to act. It reaches only what your account can, and only what the connector allows.

START TO END, IN EACH

CLAUDE

- 1 Customize > Connectors, then Discover.
- 2 Find the service; check its Verified or Community label.
- 3 Select “Connect to Claude” (on Team or Enterprise, “Request” until an owner adds it).
- 4 Sign in, read the access it asks for, approve.
- 5 In a chat: “+” > Connectors, switch it on.
- 6 When it asks to use a tool: “Allow once”, until you trust it.

CHATGPT

- 1 Settings > Plugins, then the app, or the plugin that holds it.
- 2 “Install plugin”, then “Connect”; sign in to the right account.
- 3 Review the permissions, then finish the sign-in.
- 4 In a chat, type @ or select “+” to pick the app.
- 5 While you learn: “Always ask” or “Allow read actions”.

GEMINI

- 1 Left menu > Customize, to find and connect apps (over 40, mostly US and English).
- 2 For Gmail, Drive, Docs and Calendar: the Google Workspace app. At work, your admin decides.
- 3 Mention an app with @ to make Gemini use it.
- 4 Remove apps you no longer use from Connected Apps.

BEFORE YOU CONNECT ANYTHING

- | | |
|---|--|
| ☐ Is it allowed? Use what your company approves. | ☐ Who made it? Prefer the official or verified one. |
| ☐ What can it change? Read the access screen first. | ☐ Start careful. Needs approval, or Always ask. |
| ☐ The right account. Work data, work account. | ☐ Only what it needs. Disconnect when done. |
| ☐ Hidden instructions. Outside pages can steer it. | ☐ Where data goes. Each service keeps its own terms. |

YOUR OWN MCP LINK

Each can also take an MCP server’s address directly (in Claude, Add, then Custom). Leave that to whoever runs your tools: Anthropic says custom connectors “aren’t verified”.

One small job a day A SUGGESTION

Fifteen to thirty minutes a day, using work you'd do anyway. Do it in your main assistant first, and repeat it in a second one only when you're comfortable.

DAY 1	Set up	Sign in with your company account and leave the model on its default. Write standing instructions about your role (page 5), and check what memory is on. DONE WHEN It knows who you are, and you know what it keeps.
DAY 2	Make a project	Pick one job you repeat, such as a weekly update. Make a project (Claude or ChatGPT) or a notebook (Gemini), with three to five files and a short instruction. DONE WHEN Your repeat job has a home.
DAY 3	Brief it	In that project, write a six-part brief for one real task and ask for a file. Give the brief to a colleague for the ten second test first. DONE WHEN They could follow it without asking you anything.
DAY 4	Check it	Run the five checks on what came back. Trace two names, numbers or dates outside the chat. DONE WHEN You know what was checked, and what wasn't.
DAY 5	Connect one app	Pick one connector your company allows, such as your calendar. Run the checks on page 11, keep approvals on, and ask one read-only question. DONE WHEN It answered, and changed nothing.
DAY 6	Hand over a task	Claude: ask in an ordinary conversation. ChatGPT: switch to Work. Gemini, personal account: click Task. Run a starter brief (pages 13 and 14) at its level, and read the steps it took. DONE WHEN It ran clean. That's run one of five.
DAY 7	Save a skill	Write Day 3's steps as a skill, with a plain description that says when to use it. Test it once, and fix what it misses. DONE WHEN It runs the same way every time.

TRACK FIVE CLEAN RUNS ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

Tick a box each time a brief runs clean, with nothing to fix. Our rule of thumb: five in a row, up one level. A surprise: down at once.

RUN IT IN YOUR NEXT TEAM MEETING: 15 MINUTES

0 TO 5 MIN

What changed (page 3). Ask the room: which of our tools can now act, not just answer?

5 TO 10 MIN

Each person names one job, picks its assistant and its tool (pages 4 and 6) and asks the three questions.

10 TO 15 MIN

Agree one line from Pin this up (page 15) as a team rule, and who checks.

The Team workshop runs it live, in half a day (page 18). neosaipien.com/services/#workshops

Copy one, fill the [brackets], set its level in the tool, then run its check. All twelve are also in starter-briefs.md.

Managers

01 Plan a meeting that ends in a decision	L1 READ	02 Hand over a task	L2 DRAFT
Goal: Draft an agenda for a [length] minute meeting on [date] that ends in a decision on [the decision]. Done looks like: For [who's in the room, and what each cares about]: a one-line purpose, three timed items, the exact question to answer, and what to read first. Context: [The options and any numbers], and [the notes from the last meeting on this]. May touch, must not: Use only what I paste, and reply here. Don't send invites or contact anyone. Limits: One page. Stop and ask: If the people in the room can't make this decision, say so first.		Goal: Write a handover for [the task], which [first name] takes over on [day]. Done looks like: For [first name], who knows [what they know] but hasn't done [the new part]: what done looks like, the deadline, who to ask, what they can decide alone, and two things that could go wrong. Under 200 words. Context: [My notes, links and a past example of the finished task]. May touch, must not: Read those files and save one draft. Don't share it or message anyone. Limits: One draft. Stop and ask: If my notes leave out a deadline or an owner, ask me rather than guess.	
CHECK The job. The minutes add up to the meeting's length, and this room can answer the question.		CHECK The job. Read it as them on day one. Could you start without asking anything?	

NEVER PASTE Performance ratings, pay, health details or complaints about anyone on your team.

Sales

03 Prepare for objections	L1 READ	04 Reopen a deal that went quiet	L2 DRAFT
Goal: Prepare me for the objections I may hear about [the offer] on [date]. Done looks like: For me, meeting [the buyer: a cautious finance director]: five likely objections, each with what might sit behind it, an honest two-sentence answer and a question to ask back. Context: The facts I can stand behind: [terms, timelines, what's included]. May touch, must not: Use only those facts, and reply here. Don't search the web or open my email. Limits: Half a page. Stop and ask: If an answer needs a fact I haven't given you, flag it instead of filling it in.		Goal: Draft a short email to [first name], who went quiet after [the last step: our proposal], to get a clear answer, even if it's no. Done looks like: Under 80 words, no guilt and no checking-in opener, ending with one easy question they can answer in a line. Context: Our last exchange: [paste it, or point to the thread]. They're probably busy with [what you know]. May touch, must not: Read that thread and save one draft in my Drafts. Never send it. Limits: One draft. Stop and ask: If the thread shows they already said no, tell me instead of drafting.	
CHECK The facts. Every answer rests on the facts you gave it. Cut any promise or figure you can't back up.		CHECK The sign-off. The tone is one you'd use face to face, and you press send yourself.	

NEVER PASTE Another customer's prices or terms, or contact details your rules don't allow.

Operations

05 Write up what went wrong	L1 READ	06 Plan a rollout	L2 DRAFT
Goal: Help me write up [the incident: the late delivery on Tuesday], so we agree the cause and change [the one thing]. Done looks like: For [the operations team and our director], who want facts, not blame: what happened in time order, likely causes, what we'll change with an owner and a date, and what we still can't answer. Context: The timeline and messages: [paste, with personal details taken out]. May touch, must not: Use only what I paste, and reply here. Don't name or blame a person. Limits: One page. Stop and ask: If two messages give different times, show both and ask me.		Goal: Draft a rollout plan for [the change: a new booking system], live on [date] without breaking [orders, payroll]. Done looks like: For [the people affected: 30 staff on two sites, some on nights]: phases with dates, who does what, how we tell people, the checks before each step, and how we go back if it fails. Context: [Dates, limits and who's involved], and [the calendar of holidays and busy weeks]. May touch, must not: Read those files and save one draft. Don't book, invite or message anyone. Limits: Two pages at most. Stop and ask: If a date clashes with a holiday or a busy week, flag it rather than move it.	
CHECK The facts. Every time and event matches your records. Treat the causes as ideas to test, not findings.		CHECK The changes. One new draft, nothing booked or sent. Then give each task an owner who agrees.	

NEVER PASTE Door or alarm codes, system passwords, or staff home addresses.

HR People's details stay out: a first name or a role is enough.

07 Explain a policy to a new starter	L2 DRAFT
Goal: Rewrite our policy on [the topic: expenses] in plain words for [first name, or the new starter], who joins on [date]. Done looks like: What they can do, what they need to do and who to ask, in under 200 words, with the steps as a list and a closing line saying the full policy is the one that counts. Context: The policy text: [paste, or the policy file]. May touch, must not: Read the policy and save one draft. Don't send it or open staff records. Limits: One draft, in about ten minutes. Stop and ask: If the policy is unclear or contradicts itself, quote both parts and ask me.	
CHECK The facts. Read it beside the original. Every limit, date and condition must survive.	

08 Find the themes in survey comments	L1 READ
Goal: Group these staff survey comments into themes, so [the leadership team] can pick [two or three] to act on. Done looks like: Five themes at most, each with a one-line summary, roughly how many comments fit and two short quotes, then the comments that fit nowhere. Context: The comments, with names and small-team details removed: [paste]. May touch, must not: Use only what I paste, and reply here. Don't guess who wrote a comment. Limits: Five themes. Stop and ask: If a comment raises a safety or wellbeing concern, flag it to me on its own.	
CHECK The facts. Spot-check the counts, and find every quote word for word in the survey.	

NEVER PASTE Health, pay, disciplinary or grievance details, or anything that shows who wrote a survey comment.

Marketing

09 Rewrite a web page for its reader	L2 DRAFT
Goal: Rewrite our page on [the service] for [the reader: someone comparing us with two others], so they see what we do and [book a call]. Done looks like: A headline under ten words, three short sections and one call to action. Every fact kept, none added. A draft copy, not published. Context: The current page: [paste, or its address]. May touch, must not: Read the page and save one draft copy. Don't edit or publish the live page. Limits: About the original's length. Stop and ask: If something looks out of date, flag it rather than fix it.	
CHECK The facts. Compare it line by line: no new facts, no lost conditions, no new promises.	

10 Plan a customer interview	L1 READ
Goal: Write questions for a [30] minute customer story interview with [their role] at [the kind of organisation]. Done looks like: Ten open questions in order, two follow-ups for the key ones, and a short note to send them first. No leading questions. Context: What we already know about their project: [paste]. May touch, must not: Use only what I paste, and reply here. Don't look the customer up or contact them. Limits: One page. Stop and ask: If the story would need a result or a figure, make it a question for them, not a claim.	
CHECK The job. Every question is open and none leads. Any result or figure is asked about, never assumed.	

NEVER PASTE Unreleased plans, customer lists, or a customer's results they haven't approved.

Finance Results, forecasts and contracts are often off limits: paste only what your rules allow.

11 Pull the money terms from a contract	L1 READ
Goal: List the payment dates, price rises, penalties and notice periods in this contract with [the other party], so [who will manage the account] knows what to track. Done looks like: A table: the term, what it says in plain words, the clause number and any date, with a clause number on every row. Say so if something isn't there. Context: The contract text: [paste, if your rules allow it]. May touch, must not: Use only this contract, and reply here. Don't fill gaps from other contracts or the web. Limits: One table, one page at most. Stop and ask: If a clause is ambiguous, quote it and ask rather than interpret it.	
CHECK The facts. Open each clause number and read the original. The table is a map, not the contract.	

12 Chase an overdue invoice	L2 DRAFT
Goal: Draft a reminder to [first name, their role] about invoice [number], so it's paid by [date] and the relationship holds. Done looks like: Under 100 words, polite and firm: the invoice number, amount, original due date and who to call. Leave a [how to pay] gap for me. Context: The invoice details: [paste, if your rules allow it], and [our credit policy for this stage]. May touch, must not: Read the invoice and the thread, and save one draft. Never send it or change the account. Limits: One draft. Stop and ask: If the thread shows a dispute or a promised payment, tell me instead of drafting.	
CHECK The facts. The number, amount and dates are copied exactly, and the tone fits your credit policy.	

NEVER PASTE Bank or card details, payment links, or results that aren't public yet.

Every brief is an illustration of a common task and names no real company or person. L1 Read: it reads and replies here. L2 Draft: it saves a draft only you can see. Check your workplace's rules before you paste.

Stay in charge of the work

BEFORE YOU HAND IT OVER

- ☐ It has a goal, and says what done looks like.
- ☐ It says what it may touch, and what it must not.
- ☐ It says when to stop and ask.
- ☐ It passes the ten second test.
- ☐ Its level is set in the tool, not only typed.

BEFORE YOU PASS IT ON

- ☐ **The job:** it did what the brief asked.
- ☐ **The facts:** names, numbers, dates and quotes traced outside the chat.
- ☐ **The steps:** you read how it got there.
- ☐ **The changes:** you looked where it acted.
- ☐ **The sign-off:** you'd put your name to it.

NEVER

- Paste passwords, security codes or payment details.
- Take "I checked it" as checked. Ask what was checked.
- Let an agent that reads outside content send, pay or delete without asking.
- Schedule a task that opens sensitive files, sends messages for you or buys things.

THREE QUESTIONS SET THE CEILING

- Can it be undone easily?**
- Does it leave the team?**
- Does it read outside content?**

FOUR LEVELS, SET IN THE TOOL

- 1 **Read**
- 2 **Draft**
- 3 **Ask first**
- 4 **Alone**

Our rule of thumb: up one level after five clean runs. Down at once after a surprise.

"The AI did it" is never the answer to "who checked this?"

Where every fact comes from

Every product fact here is as of October 2026. Each maker's page below was read on 16 Oct 2026; a date in a row is the day the maker published it. The research behind the habits is on page 17.

CLAUDE (ANTHROPIC)

Page 3	Introducing Claude Opus 5.5, 22 Sep 2026. anthropic.com/news/claude-opus-5-5
Page 3	Models overview (Claude Fable 5.1, 1 Sep 2026). platform.claude.com/docs/en/about-claude/models/overview
Pages 3, 5, 12	Claude Cowork and chat are now one Claude, 16 Sep 2026. claude.com/resources/articles/cowork-is-now-claude
Page 4	Dashboards and motion; Docs, Slides and Design, 8 Oct 2026. claude.com/resources/articles/dashboards-and-motion
Pages 4, 5	Claude Cowork product page and FAQ. claude.com/product/cowork
Page 5	Claude Code product page. claude.com/product/claude-code
Page 4	Claude plans and pricing. claude.com/pricing
Page 4	Use Claude for Excel. claude.com/docs/office-agents/excel.md
Page 5	What are projects. support.claude.com/en/articles/9517075-what-are-projects
Page 5	Understanding Claude's personalization features. support.claude.com/en/articles/10185728-understanding-claude-s-personalization-features
Page 10	Use Claude's chat search and memory. support.claude.com/en/articles/11817273-use-claude-s-chat-search-and-memory-to-build-on-previous-context
Page 5	Change the model, effort, and thinking settings. support.claude.com/en/articles/8664678-change-the-model-effort-and-thinking-settings
Page 5	Choosing the right Claude model. academy.claude.com/tutorials/choosing-the-right-claude-model
Page 5	Use research on Claude. support.claude.com/en/articles/11088861-using-the-research-feature
Pages 3, 5, 8	Skills overview. claude.com/docs/skills/overview.md
Page 11	Connectors, skills, and plugins. claude.com/docs/extend/overview
Pages 5, 10, 11	Get started with connectors. claude.com/docs/connectors/overview
Page 11	Use the connectors directory. support.claude.com/en/articles/11724452-use-the-connectors-directory-to-extend-claude-s-capabilities
Pages 10, 15	Use Claude Cowork safely. support.claude.com/en/articles/13364135-use-claude-cowork-safely
Page 3	Effective context engineering for AI agents, 29 Sep 2025. anthropic.com/engineering/effective-context-engineering-for-ai-agents
Page 7	Prompting best practices (the new employee, the colleague test, examples). platform.claude.com/docs/en/build-with-claude/prompt-engineering/claude-prompting-best-practices

CHATGPT (OPENAI)

Pages 3 to 5	GPT-6 Astra, 3 Sep 2026. openai.com/index/gpt-6-astra/
Page 5	Introducing GPT-6.1 Sol, 29 Sep 2026. openai.com/index/introducing-gpt-6-1-sol/

Pages 3, 4, 17	ChatGPT release notes (9 Jul, 11 Sep and 29 Sep 2026). help.openai.com/en/articles/6825453-chatgpt-release-notes
Pages 3 to 5, 12	ChatGPT Work and Codex. help.openai.com/en/articles/20001275-chatgpt-work-and-codex
Page 5	GPT-6 and other models in ChatGPT. help.openai.com/en/articles/20001354-gpt-6-and-other-models-in-chatgpt
Page 4	ChatGPT plans and pricing. chatgpt.com/pricing/
Page 5	Projects in ChatGPT. help.openai.com/en/articles/10169521-projects-in-chatgpt
Page 5	ChatGPT custom instructions. help.openai.com/en/articles/8096356-chatgpt-custom-instructions
Page 10	Memory FAQ. help.openai.com/en/articles/8590148-memory-faq
Pages 10, 11	Connected apps in ChatGPT. help.openai.com/en/articles/11487775-apps-in-chatgpt
Pages 5, 11	Skills and plugins. learn.chatgpt.com/docs/skills-and-plugins.md
Pages 4 to 6	Deep research. help.openai.com/en/articles/10500283-deep-research
Pages 4, 10, 15	Using cloud browser in ChatGPT. help.openai.com/en/articles/20001280-using-cloud-browser-in-chatgpt

GEMINI (GOOGLE)

Pages 4, 5, 10	Gemini Apps release notes. gemini.google/release-notes/
Page 3	Introducing Gemini 4 Argon, 13 Oct 2026. blog.google/innovation-and-ai/models-and-research/gemini-models/introducing-gemini-4-argon/
Page 5	About Gemini models. support.google.com/gemini/answer/18648722
Pages 3, 5, 12, 17	Manage your tasks and workflows in Gemini Apps. support.google.com/gemini/answer/17094507
Pages 5, 8	Create and manage skills for Gemini Apps. support.google.com/gemini/answer/17094296
Pages 3, 17	About the transition from Gems to skills. support.google.com/gemini/answer/18560919
Pages 10, 11	Connect and manage custom apps for Gemini Apps. support.google.com/gemini/answer/17209137
Page 11	Use apps connected to Gemini with a work or school account. support.google.com/gemini/answer/14959807
Page 4	Get started with Google Workspace with Gemini. support.google.com/mail/answer/13952129
Pages 4, 5, 17	NotebookLM is now Gemini Notebook, 16 Jul 2026. blog.google/innovation-and-ai/products/gemini-notebook/notebooklm-gemini-notebook/
Pages 4, 5	Use Deep Research in Gemini Apps. support.google.com/gemini/answer/15719111
Page 4	Google AI plans. gemini.google/subscriptions/

THE TWO SHARED STANDARDS

Page 11	What is the Model Context Protocol. modelcontextprotocol.io/docs/getting-started/intro
Pages 3, 8	Agent Skills overview. agentskills.io/

The research behind the habits

The research and the studies behind the three habits. Each page was read on 16 Oct 2026 unless the row says otherwise; each row names the primary sources its wording was checked against.

REGULATORS, RESEARCH AND SURVEYS

Page 3	UK National Cyber Security Centre, Managing the cyber risk of agentic AI, 20 Aug 2026. nscs.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai
Page 6	Gartner, press release, 25 Jun 2025. gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027
Page 9	Microsoft, 2026 Work Trend Index, 5 May 2026. microsoft.com/en-us/worklab/work-trend-index/agents-human-agency-and-the-opportunity-for-every-organization
Page 9	University of Melbourne, media release on the global study 2025. fbe.unimelb.edu.au/newsroom/media-release-global-study-reveals-trust-of-ai-remains-a-critical-challenge-reflecting-tension-between-benefits-and-risks
Page 9	KPMG, Trust, attitudes and use of AI: a global study 2025. kpmg.com/xx/en/our-insights/ai-and-technology/trust-attitudes-and-use-of-ai.html
Page 3	Wharton Generative AI Labs, Prompting Science Report 3, Meincke, Mollick, Mollick and Shapiro, 1 Aug 2025 (checked 6 Oct 2026). arxiv.org/abs/2508.00614
Page 7	Elizabeth Louise Newton, PhD dissertation, Department of Psychology, Stanford University, June 1990 (checked 14 Oct 2026)

WHERE THE WORDING COMES FROM

Page 9	One word at a time; likely isn't the same as true: OpenAI, Why language models hallucinate, 5 Sep 2025; Anthropic, Tracing the thoughts of a large language model, 27 Mar 2025 (checked 4 and 7 Oct 2026)
Pages 9, 15	"Are you sure?" isn't checking; quotes and links can be invented too: Mata v. Avianca, US District Court, Southern District of New York, 22 Jun 2023; Sharma and others, arXiv 2310.13548; Laban and others, arXiv 2311.08596; Columbia Journalism Review, Tow Center, 6 Mar 2025 (checked 4 and 7 Oct 2026)
Pages 9, 15	Ask what was checked: Bastounis, Circelli and Hansen, preprint, arXiv 2610.08144 (checked 14 Oct 2026)
Pages 6, 10	Workflow or agent; the longer leash: Anthropic, Building effective agents, Dec 2024; UK NCSC, 20 Aug 2026 (checked 3 Oct 2026)
Page 8	Six words leave the AI to guess; it usually just answers with a safe, typical choice: the makers' own prompting guides (checked 6 Oct 2026)
Pages 10, 15	Outside content and prompt injection; these steps cut the risk: UK NCSC, 10 Dec 2025; OWASP, LLM01:2025 Prompt Injection (checked 3 Oct 2026)

IF A BUTTON ISN'T THERE

Features reach people in waves. If you can't see something in this guide, the usual reasons are a rollout still on its way, your plan, your country, or a setting your admin controls. These products change weekly: before you rely on a menu path, check the date here and your tool's own help page.

NAMES ON THEIR WAY OUT

The ChatGPT Atlas browser stopped working on 9 Aug 2026; its features moved into ChatGPT and Codex. Custom GPTs are moving to plugins. Gems become skills from Nov 2026 for personal accounts, and in 2027 for work and school ones. NotebookLM is now Gemini Notebook, and Gemini Spark is now tasks in Gemini.

Nothing in this guide is legal advice. The ranking on page 4 is my own view, not a maker's claim.

Keep learning, free

The Field Guide, level 1, Start here. neosaipien.com/learn/

Every kit, free

The library: a team AI policy, a fact-check kit, an agent playbook, two workshops and a 90-day rollout plan, all under the same free licence. neosaipien.com/library.html

BRING THIS TO YOUR TEAM

I'm Samson. Neosaipien is my AI training and consulting practice, and I run these live, in the Team workshop.

neosaipien.com/services/

BOOK A CALL

neosaipien.com/book.html

WAYS TO WORK WITH ME

Leadership briefing

90 minutes with one leadership team: the real risks, and the decisions to make.

Executive coaching

Six one-to-one sessions on a leader's own work, until AI is part of how they do it.

Team workshop

Half a day of hands-on training built around your team's real work.

Workflow automation

One sprint: repetitive work mapped, automated and handed back.

AI implementation

A short discovery first, then AI put to work inside the tools you already run.

ABOUT

Made by Neosaipien, AI training and consulting for the people doing the work: the Leadership briefing, Executive coaching, the Team workshop, Workflow automation and AI implementation. Every fact in this guide was checked against its source (pages 16 and 17).

LICENCE

Free from Neosaipien under the free licence. Use it, change it for your team and share it free, keeping the credit and the date. A changed version stays inside your own organisation. Please don't sell it. Full terms: LICENCE.txt, in the download.

EDITION

October 2026 edition. Checked on 16 Oct 2026. These products change monthly: check the settings in your own tool. Not legal advice.

neosaipien

Neosaipien is not affiliated with Anthropic, OpenAI or Google.

neosaipien.com
samson@neosaipien.com