

[◆ FREE GUIDE](#)

AI ESSENTIALS · GUIDE NO. 21

WAKE UP TO A FINISHED INDUSTRY BRIEF

A standing AI job that reads a fixed list of sources every morning, checks each item against yesterday's page so the same story stops reappearing, separates what is confirmed from what is only claimed, and writes the whole thing into Notion before you sit down.

~12 MIN READ · FOR EVERYONE WHO COMMENTED "MONITOR" · FIRST SETUP: ~30 MINUTES ·
CAPABILITIES VERIFIED JULY 2026

Hank Barker // AI Essentials // @hankhelpsai

THE SHORT VERSION

- The quality of the brief is decided by the **source list**, not the model. Fix the list first: first-party sources, trusted publications, named people.
- Every item carries **five fields**: what changed, when it happened, the original source, who it affects, what action may be worth taking. Anything missing one of them is noise.
- Each run **reads yesterday's page first**, so the same story stops showing up for a week.
- Two tiers of trust: **confirmed** when it traces to a primary source, **unverified** when it does not. Plus an **Ignore for now** section, because removing noise is half the value.
- Built as a Claude Cowork scheduled task writing to Notion. A ChatGPT version is at the end.

WHY "ASK AI WHAT'S NEW" NEVER WORKS

A summary of whatever the model happened to find is not a brief. A brief is a promise that the same ground gets covered the same way every single morning.

Most professionals have tried some version of this. You open a chat, ask what changed in your industry this week, and get back something fluent and plausible. It reads well. It is also unrepeatable. Ask the same question twice and you get two different lists, in a different order, with different things missing. There is no way to tell which items are real, which are one blog post repeating another blog post, and which are simply what surfaced first.

That unreliability is the actual problem, and a better prompt does not fix it. What fixes it is removing the model's discretion over what counts. Fix the sources, fix the required fields, and make yesterday's output an input to today's run, and the brief becomes something you can trust at 7 AM without re-checking it yourself.

This guide builds that. Five parts, in order: the source list, the five required fields, the de-duplication pass, the confirmed and unverified split, and a section for what got deliberately ignored. Then you schedule it and point the output at a Notion page. The parts are simple. Holding all five at once is what makes the result different from anything you have tried before.

**You are not asking the AI to be smart about your industry.
You are asking it to be relentlessly consistent about a list
you wrote.**

STEP 1 · BUILD THE SOURCE LIST

This is the step people skip, and it is the step that decides whether any of the rest matters. An hour spent here is worth more than every prompt tweak you will ever make. Your source list has three tiers, and they do different jobs.

THE SOURCE LIST IS THE WHOLE BRIEF



Tier 1 First-party sources

The organization that actually made the change. Filings, changelogs, official notices, standards bodies.



Tier 2 Trusted publications

The two or three outlets that cover your corner properly and get corrected when they are wrong.



Tier 3 Specific people

Named practitioners whose posts move before the trade press catches up. People, not accounts you

Three tiers, three jobs. First-party sources are the evidence. Publications are the interpretation. People are the early warning.

Tier 1 · First-party sources

These are the organizations that make the change, publishing it themselves. Regulator notices and filings. Official product changelogs and release notes. Standards bodies. Company newsrooms and investor pages. A court docket, if that is your world. The defining quality is that nobody is between you and the fact.

Aim for six to ten of these. They are the backbone, and they are the only sources that can put an item in your Confirmed section. If you cannot name six first-party sources for your own field, that gap is itself a finding, and closing it is the most valuable hour in this whole setup.

Tier 2 · Trusted publications

Two or three outlets that cover your corner properly. The test is not readership, it is whether they issue corrections and whether their reporters name their sources. A trade newsletter run by one well-connected person often beats a large publication that covers your field twice a year.

Publications do a job first-party sources cannot: they tell you what a change means and who is upset about it. That interpretation is genuinely useful. It is also, by definition, secondary, which is why it lands in a different section of your brief.

Tier 3 · The specific people

Name individuals, not vague topic areas. The compliance lead who posts about draft rules before the trade press picks them up. The engineer at your main vendor who explains what actually shipped. The competitor's head of product. Five to eight people is plenty, and they are usually your earliest signal by a day or two.

Be honest about who is on this list. The test is whether that person's post has ever changed something you did. If not, they are entertainment, and entertainment belongs in your feed, not in your brief.

Assembling it without overthinking

The fastest way to build the list is backwards. Take the last five things that genuinely mattered in your work: a rule that changed, a competitor move, a vendor decision, a pricing shift, a deadline. For each one, write down where you first heard about it and where the original evidence lived. Forty minutes of that usually produces twelve to fifteen sources, and they are demonstrably the right ones, because they already proved it.

Keep the total between twelve and twenty. Below twelve you will miss things. Above twenty the brief gets long and you stop reading it. Review the list once a quarter: add anything that surprised you, cut anything that has never produced a useful item.

RSS, and why a feed beats "go read the internet"

Most publications, newsletters, blogs, and a good number of official newsrooms still publish an RSS feed, which is just a stable address listing recent posts with their dates and links. Collecting those feed addresses into your source list is a reasonable way to hand the AI something fixed and reliable to work from, rather than a topic and a hope.

Be honest about the difference. A curated feed you chose is dependable, because the address does not move and the dates are structured. Asking an AI to go find whatever it can on arbitrary websites is far less reliable: pages change, some sites cannot be read, and what comes back varies run to run. Where a feed exists, use it. Where one does not, name the specific page you want checked, and accept that a login-walled site may need you to check it yourself.

THE CORE IDEA OF THIS WHOLE GUIDE

The AI's own web access is not a substitute for your source list. It is a way to reach the list. If the model finds something outside your sources, that is not a bonus, it is an unverified claim until it traces back to a primary document. Anything it cannot trace gets labeled, not deleted and not promoted.

STEP 2 - THE FIVE FIELDS EVERY ITEM CARRIES

An item without these five fields is a headline, and a page of headlines is exactly the thing you were trying to escape. Requiring all five does something subtle and important: it makes the AI unable to include anything it does not actually understand, because it cannot fill in "who it affects" for a story it only half-read.

ANATOMY OF ONE WELL-FORMED ITEM

Regulator updates the quarterly reporting template

CONFIRMED

- 1

WHAT CHANGED
The template gained two required fields. One sentence, specific, no adjectives.
- 2

WHEN IT HAPPENED
The date on the notice itself, not the date something wrote about it.
- 3

THE ORIGINAL SOURCE
A link to the notice. Not a link to coverage of the notice. This field is the whole trust model.
- 4

WHO IT AFFECTS
Named accounts, teams, or clients of yours. Not "the industry".
- 5

WHAT ACTION MAY BE WORTH TAKING
One concrete next step, or the words "no action yet".

One item, five fields. Field 3 is doing the heaviest lifting: it decides which section the item lands in.

FIELD	WEAK VERSION	WHAT YOU WANT INSTEAD
What changed	"Big update to reporting rules."	"The quarterly template gained two required fields." One sentence, no adjectives.
When it happened	"Recently."	The date and time on the source document itself, not the date it was written about.
The original source	A link to an article about it.	A direct link to the notice, filing, or changelog. If none exists, say so.
FIELD	WEAK VERSION	WHAT YOU WANT INSTEAD
Who it affects	"The industry."	Named clients, accounts, or teams of yours. This is the field only you can define well.
What action may be worth taking	"Worth monitoring."	One concrete step, or the words "no action yet". Both are useful. Vagueness is not.

The fourth field is the one that makes the brief yours rather than generic. Name your clients, accounts, portfolio companies, or internal teams in the instruction, and the AI can say "this lands on the three accounts filing on the quarterly cycle" instead of "this affects many organizations". That one sentence is the difference between a page you skim and a page that changes your morning.

LET IT SAY "NO ACTION YET"

If every item demands a next step, you will get invented ones. Explicitly permitting "no action yet" keeps the action field honest, and it makes the two or three items that do carry a real action stand out on the page.

STEP 3 - COMPARE AGAINST YESTERDAY'S PAGE

Here is what kills most automated briefs by week two. A story breaks on Monday. It gets covered Monday afternoon, analyzed Tuesday, aggregated Wednesday, and referenced in a roundup Thursday. Without a de-duplication step, your brief reports it four times, you learn to skim, and the whole thing quietly becomes wallpaper.

The fix is to make yesterday's output an input. Before writing anything, the run opens the previous page and reads it. Anything already there is dropped. Anything that materially advanced gets one new line under the existing item, saying what is new and nothing else. Only genuinely new items get written up in full with all five fields.

READ YESTERDAY FIRST, THEN WRITE TODAY

YESTERDAY'S PAGE

Reporting template change
Vendor pricing revision
Draft rule out for comment

COMPARE

TODAY'S PAGE

~~Reporting template change~~
~~DROPPED, ALREADY THERE~~
Vendor pricing revision
ONE NEW LINE ADDED
Certification deadline moved
NEW TODAY

Three outcomes per story: already covered and dropped, advanced and given one line, or new and written up in full.

Two details make this work. Put the actual address of yesterday's page in the instruction rather than trusting the AI to find the right one, and define what "materially advanced" means. A draft becoming final is material. A third outlet writing about the same draft is not.

A USEFUL SIDE EFFECT

Once the run reads yesterday before writing today, your Notion page becomes a running record of when each change actually landed. Three weeks in you can scroll back and reconstruct the timeline of a story, which is something no feed reader has ever given you.

STEP 4 - CONFIRMED, OR UNVERIFIED

Every item goes in one of two trust tiers, and the rule is mechanical rather than a judgement call. If field three points at a primary document you can open, the item is Confirmed. If it does not, the item is Unverified, and it carries a one-line note saying what is missing.

The important half of that rule is what does not promote an item. Repetition does not. Five outlets carrying the same story, all citing each other, is one unverified claim wearing five hats. Confidence does not either: a well-written, self-assured summary is not evidence. Only a document you can open moves an item into Confirmed.

EVERY ITEM LANDS IN EXACTLY ONE OF THREE SECTIONS

Confirmed

Field 3 points at a primary document.

Official notice
Published changelog
A filing you can open
A named person, on the record, first-hand

Unverified

Probably real. Cannot be traced to a source.

"People familiar with"
Outlet quoting outlet
A screenshot, no link
Anything the AI could not open and read

Ignore for now

Looked at, deliberately left out, with a reason.

Real, but not yours
Announced, not shipped
Recycled from last month
Too early to act on

One mechanical rule per section. Nothing is deleted for being unverified, it is labeled, which is what lets you act on the top section without checking it.

This is not about being pedantic. It is what makes the brief usable at speed. When you trust that Confirmed means traceable, you can act on that section directly and treat Unverified as a watch list. Without the split you have to re-check everything yourself, which is exactly the work you were trying to hand off.

KEEP UNVERIFIED ITEMS VISIBLE

The temptation is to have the AI drop anything it cannot verify. Resist it. A rumour about your biggest vendor is worth knowing about even when nobody will confirm it. Labeled, it is useful intelligence. Deleted, it is a blind spot you do not know you have.

STEP 5 - THE IGNORE FOR NOW SECTION

A brief that only tells you what to look at is doing half the job. The other half is telling you what you can safely stop thinking about, and that half is what most people never build.

So the page gets a third section listing what the run looked at and deliberately left out, one line each, with the reason. Real but not relevant to your accounts. Announced but not shipped. A recycled version of something from last month. Too early to act on. Interesting but outside the scope you defined.

Two things happen when you add this section. A story you half-remember shows up with a reason attached, which closes the loop instead of leaving it nagging. And it becomes your tuning instrument: every time you disagree with something sitting in Ignore for now, that is a precise correction to make to the standing instruction. Without it you only ever see what the AI included, never what it quietly threw away.

Anything the run touched and discarded should be visible. What gets thrown away tells you more about whether the system is working than what gets kept.

THE STANDING PROMPT, PART ONE

Here is the instruction that encodes all five parts. Replace everything in brackets with your own sources, accounts, and page addresses. Part one sets the sources, the window, and the required fields.

You are my industry monitor. Run this every weekday morning.

SOURCES. Use only these, in this order: [first-party source 1], [first-party source 2], [first-party source 3], [publication 1], [publication 2], [person 1], [person 2]. Do not add sources on your own. If one is unreachable, note it at the bottom of the page rather than substituting another.

WINDOW. Cover only items published in the last 24 hours. On Monday, cover the last 72 hours.

FIVE REQUIRED FIELDS. Do not include an item that is missing any of these:

1. What changed. One sentence, specific, no adjectives.
2. When it happened. The date and time on the source itself.
3. The original source. A direct link to the primary document, not to coverage of it.
4. Who it affects. Which of these is touched: [account or team A], [account or team B], [account or team C]. If none, say "none of mine".
5. What action may be worth taking. One concrete step, or the words "no action yet".

THE STANDING PROMPT, PART TWO

Part two is what keeps the brief honest over time: the de-duplication rule, the trust split, the ignore section, and where the finished page goes. Paste both parts as one instruction.

DE-DUPLICATION. Before writing anything, open yesterday's brief at [Notion page URL] and read it. If a story already appears there, do not repeat it. If it materially advanced, meaning a draft became final or a date or number changed, add one line under the existing item saying what is new, and nothing else.

TRUST. Put an item under "Confirmed" only if field 3 points at a primary document you were able to open. Everything else goes under "Unverified" with a one-line note on what is missing. Several outlets repeating the same claim does not make it confirmed.

IGNORE FOR NOW. Add a third section listing what you looked at and deliberately left out, one line each, with the reason. Treat this section as required, not optional.

OUTPUT. Write the finished brief to a new Notion page under [parent page], titled "Industry Brief, [date]". Three sections in this order: Confirmed, Unverified, Ignore for now. If a section is empty, write "Nothing today" rather than padding it. Keep the whole page to one screen.

KEEP THE SOURCE LIST IN THE PROMPT

Put the sources in the instruction itself or on a Notion page, not in a file on your laptop. A scheduled task that depends on a local file will only run when your machine is on, which defeats the point of it running before you wake up.

WIRING IT UP IN CLAUDE COWORK

Cowork is the agentic mode in Claude, as of July 2026: you describe an outcome, walk away, and come back to finished work. It is on Claude Desktop for macOS and Windows, on the web at claude.ai, and on mobile, and it needs a paid plan (Pro, Max, Team, or Enterprise). To start one, open Claude and choose Cowork rather than Chat in the message box.

support.claude.com/en/articles/13345190-get-started-with-claude-cowork



API DocsRelease NotesHow to Get SupportEnglish

Q Search for articles...

All Collections > Claude Cowork > Get started with Claude Cowork

Get started with Claude Cowork

Updated over 2 weeks ago

This article explains how to use [Claude Cowork](#), which brings Claude Code's agentic capabilities to knowledge work beyond coding.

Availability

Claude Cowork is available for paid plans (Pro, Max, Team, Enterprise) on:

- **Claude Desktop for macOS** — [Click here](#) to download
- **Claude Desktop for Windows** — Cowork requires the latest version of Claude for Windows. Download or update at [claude.com/download](#).
- **Web**, at [claude.ai](#)
- **Claude Mobile**, in the latest version of Claude for iOS and Claude for Android

Chat and Cowork share one home, so you start both from the same place. On any surface, find the message box and select "Cowork," then describe your task. To go back to a regular conversation, select "Chat."

Claude Cowork is in beta on web and mobile, and rolling out over the next several weeks starting with the Max plan, with more plans to follow.

Availability

What is Claude Cowork?

How Claude Cowork runs your tasks

Get started

Start a Cowork session

What to expect during a task

Choose how Claude checks with you

Add global and folder instructions

Claude Cowork plugins

Schedule recurring tasks

Usage limits

Permissions and security

Example use cases

Current limitations

Troubleshooting

Cowork is the agentic mode: describe the outcome, leave, come back to a finished page. This is the mode the whole brief runs in. — **Source:** [Claude Help](#), [Get started with Claude Cowork](#)

Connect Notion

The official Notion integration lets an AI tool read and write your Notion pages the same way you can, so it can create the brief page rather than handing you text to paste. If you are not technical, connect Notion from Claude's connector directory rather than adding a raw MCP address. It is the same integration with none of the setup.

@HANKHELPSAI

AI ESSENTIALS · GUIDE NO. 21 · PAGE 10 / 15

developers.notion.com/docs/mcp

Developer Docs

Search... Ctrl K Ask Assistant Log in Developer portal >

Guides API Reference CLI Reference Workers Admin API Changelog Examples

Get started
Overview
Quickstart
Personal access tokens
Internal connections
Public connections
Authorization
Handling API keys
Preparing for users
List on the Marketplace

Agent APIs
Notion MCP
Overview
Connecting to Notion MCP
Supported tools
Security best practices
Integrating your own MCP client
Hosting a local MCP server >

API versions
Upgrading to 2026-03-11 >
Upgrading to 2025-09-03 >

Notion MCP

Learn how to connect AI agents to your Notion workspace.

Connect your AI tools to Notion using the Model Context Protocol (MCP), an open standard that lets AI assistants interact with your Notion workspace.

What is Notion MCP?

Notion MCP is our hosted server that gives AI tools secure access to your Notion workspace. It's designed to work seamlessly with popular AI assistants like Claude Code, Cursor, VS Code, ChatGPT, and more.

High-level diagram of the MCP data flow, where Notion hosts both the MCP Server and Notion! the remote

Ask a question... Ctrl+I

Notion's own developer docs for the hosted MCP server, the official route that lets an AI assistant reach your workspace. Per the same docs it reads and writes, which is what lets the brief write itself onto a page every morning. — **Source:** [Notion](#), [Notion MCP documentation](#)

READ THIS BEFORE YOU CONNECT

Authorization happens through OAuth, and it grants **full workspace access at your own permission level**. Whatever you can see in Notion, the connected tool can see, and it can write as well as read. If your workspace holds material the brief has no business touching, think about which account you connect with, or give the brief its own workspace. This is a real decision, not a checkbox.

Put it on a schedule

Cowork can run a task on a schedule. Open **Scheduled** in the left sidebar, choose **New task**, then either **Create with Claude**, which asks you multiple-choice questions and ends with a Schedule button, or **Set up manually**, where you fill in a task name, the prompt, an approval mode, and the frequency, plus an optional model and folder, then Save. Cadences available are hourly, daily, weekly, on weekdays, or manually. For a morning brief, weekdays is usually right.


[API Docs](#)
[Release Notes](#)
[How to Get Support](#)
[English](#)

 Search for articles...

All Collections > Claude Cowork > Schedule recurring tasks in Claude Cowork

Schedule recurring tasks in Claude Cowork

Updated over 2 weeks ago

Scheduled tasks allow you to delegate work to Claude Cowork by creating tasks that run automatically on a recurring basis, or on demand. Instead of starting each task from scratch, you describe it once and Claude handles it on your schedule—delivering finished outputs like reports, briefings, and summaries every time.

Scheduled tasks are available in Cowork for all paid plans (Pro, Max, Team, Enterprise).

Claude Cowork is in beta on web and mobile, and rolling out over the next several weeks starting with the Max plan, with more plans to follow.

What scheduled tasks can do

Scheduled tasks have access to the same capabilities as regular Cowork tasks, including connected tools, skills, and installed plugins. Common uses include:

- [Schedule a daily summary of Claude messages, emails, calendar events, from the past](#)

Anthropic's page for Cowork scheduled tasks. Note the box: available on all paid plans, and in beta rollout starting with the Max plan. Each run starts a fresh session with access to your connected tools, skills, and installed plugins. — **Source:** [Claude Help, Schedule recurring tasks in Claude Cowork](#)

Each run is a fresh session, which is exactly why the de-duplication step is written the way it is. The task does not remember yesterday. It reads yesterday. Your Notion page is the memory, and that is sturdier than hoping context carries over.

RUNS WHILE YOU SLEEP, WITH ONE EXCEPTION

Anthropic's wording is that scheduled tasks run remotely, so they run on their cadence even when your computer is asleep or the desktop app is closed. The exception matters: a task that requires local files or apps will only run locally. Keep your sources and your output in Notion and the prompt, and the brief stays a cloud job.

YOU MAY NOT SEE THE SCHEDULED TAB YET

As of July 2026 Anthropic states that Cowork is in beta on web and mobile, and rolling out over the next several weeks starting with the Max plan, with more plans to follow. If you are on Pro and there is no Scheduled tab, you have not done anything wrong. Build and test the prompt as a normal Cowork task now, and schedule it when the tab appears.

IF YOUR WORK LIVES IN CHATGPT

The same five-part system ports over. ChatGPT supports scheduled tasks: you open **Scheduled** in the sidebar of the desktop app, or simply ask ChatGPT in a chat to schedule a task. It handles recurring schedules, minute-based intervals, and custom RRULE schedules for anything unusual, and a task can use uploaded files, connected tools, skills, and plugins available to that chat. Note that the Codex CLI and IDE extension cannot manage scheduled tasks, so set this up in the desktop app or ChatGPT Work on the web.

The screenshot shows the 'Scheduled tasks' page in the ChatGPT Automations documentation. The page is titled 'Scheduled tasks' and has a subtitle 'Schedule recurring tasks in ChatGPT'. The main content area is divided into three sections: 'Schedule recurring tasks to run in the background', 'In the ChatGPT desktop app', and 'For project-scoped scheduled tasks'. The right sidebar contains a 'Manage scheduled tasks' section with links to 'Ask ChatGPT to create or update scheduled tasks', 'Schedule a task inside a chat', 'Test scheduled tasks', 'Worktree cleanup for scheduled tasks', 'Permissions and security model', and 'Examples'. At the bottom of the page, there is a chat interface with the prompt 'What should we work on?' and a button to 'Ask AI'.

ChatGPT's scheduled tasks are the equivalent path: recurring runs that can reach your connected tools and files. — **Source:** [OpenAI, ChatGPT automations documentation](#)

SET EXPECTATIONS CORRECTLY

Plenty of write-ups describe a separate "monitoring" trigger that only pings you when something changes. The official documentation does not define one, so build this as a recurring scheduled task and let the de-duplication step handle the quiet days. On a slow morning the page should say "Nothing today", and that is a perfectly good outcome.

THE FIRST TWO WEEKS, AND HOW TO TUNE

Expect the first week to be mediocre. That is not a failure, it is the calibration period, and knowing what to adjust makes the difference between a brief you keep and one you quietly stop opening.

- **Too many items?** Tighten the window before you cut sources. Twenty-four hours with a Monday catch-up filters more than people expect.
- **Items you do not care about?** Do not delete the source. Sharpen field four by naming your accounts and teams more precisely, and let the item fall into Ignore for now.
- **Same story reappearing?** Your definition of "materially advanced" is too loose. Say plainly that new coverage of an old fact is not an advance.
- **Everything landing in Unverified?** Your first tier is thin. Go add first-party sources, because nothing else can promote an item.
- **Ignore for now is empty?** The run is not casting wide enough, or it is not being honest about what it discarded. An empty ignore section on a busy week is a warning sign.

Set one review point at the end of week two. Read the last ten pages in a single sitting and ask one question: did this ever tell me something I would have missed, or something I could stop worrying about? If the answer is no, the problem is almost always the source list, not the prompt.

START HERE TODAY

Do not build the whole thing this morning. Open a document and write down the last five things that genuinely changed your work, and where you first saw each one. That list is your Tier 1 and Tier 3. Connect Notion, paste both halves of the standing prompt into a single Cowork task, and run it once by hand before you schedule anything. Schedule it only after a manual run produces a page you would actually read.

SOURCES & WHERE TO LEARN MORE

- [Claude Help, Get started with Claude Cowork](#)
- [Claude Help, Schedule recurring tasks in Claude Cowork](#)
- [Notion, Notion MCP documentation \(read and write, OAuth, workspace access\)](#)
- [Claude, manage your connectors](#)
- [OpenAI, ChatGPT automations and scheduled tasks](#)
- [Want the connector playbook? Comment "SYSTEM" for that guide](#)

Built by Hank Barker // AI Essentials // @hankhelpsai. Every product capability above was checked against official documentation on 31 July 2026. These tools ship fast, so re-check the menu paths if something looks different on your screen.