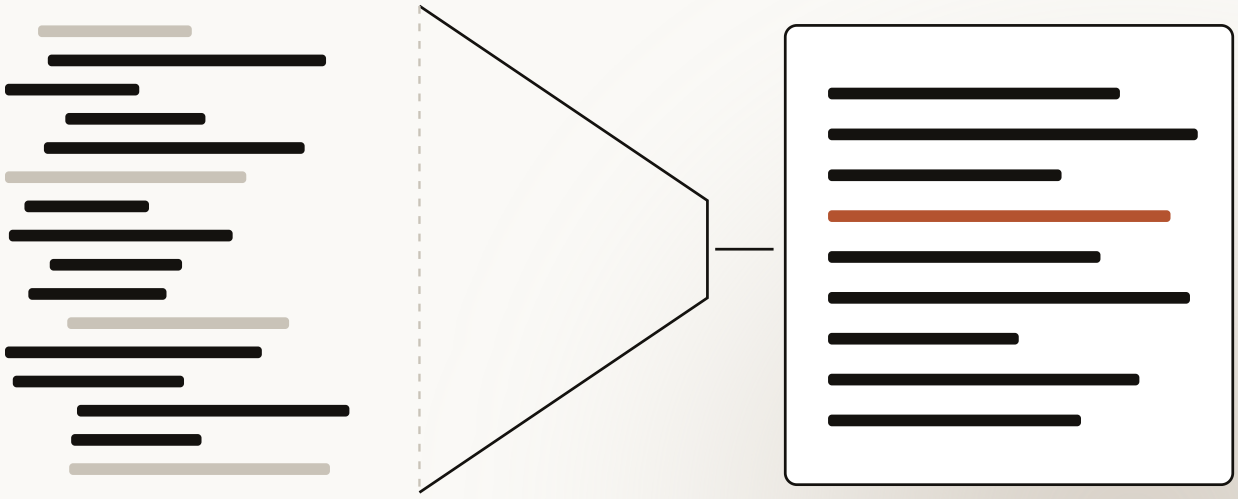


Your First AI Agent in a Weekend: a research agent that sends you a Monday briefing

Build an agent that reads your sources all week, scores every item against what matters to you, logs what it read, and emails you one checked page every Monday.

RESEARCH & WEEKLY REPORT EDITION · FREE SAMPLE

NEUROVA AI
SOFTWARE DEVELOPMENT

About this guide

GUIDE	Your First AI Agent in a Weekend: a research agent that sends you a Monday briefing
EDITION	Research & weekly report
VERSION	1.0 · last tested November 2026
PUBLISHER	Neurova AI · neurova-ai.com

Contents

01 **What a research agent is, and isn't**
Gather, judge, log, brief. And the three rules that keep a summary honest.

02 **Set up in an hour**
in the full guide

03 **Build 1: the collector**
in the full guide

04 **Build 2: the Monday briefing**
in the full guide

05 **Test it, and trust it**
in the full guide

06 **Run it safely**
in the full guide

This free sample contains chapter 01. The full guide has all 6 chapters, plus the prompt library and the checklist.

The promise. In one weekend, about five hours, you'll build two automations in Zapier: a **collector** that scores every new item from your feeds against a card you wrote and logs it, and a **Monday briefing** that turns last week's best items into one page, checks every sentence against its source, and emails it to you. Every line links to where it came from.

What you need. A Gmail mailbox, Google Sheets, and a **paid Zapier plan**. The free plan only runs two-step automations, and this build needs more. Most plans start with a free trial long enough to build and test everything here.

Licence. This is a free sample. You are welcome to share it as it is.

Examples. The worked-example business in this guide is fictional. The build is real: every step was carried out on the tools named, on the date above. Case studies from real companies are cited with their source.

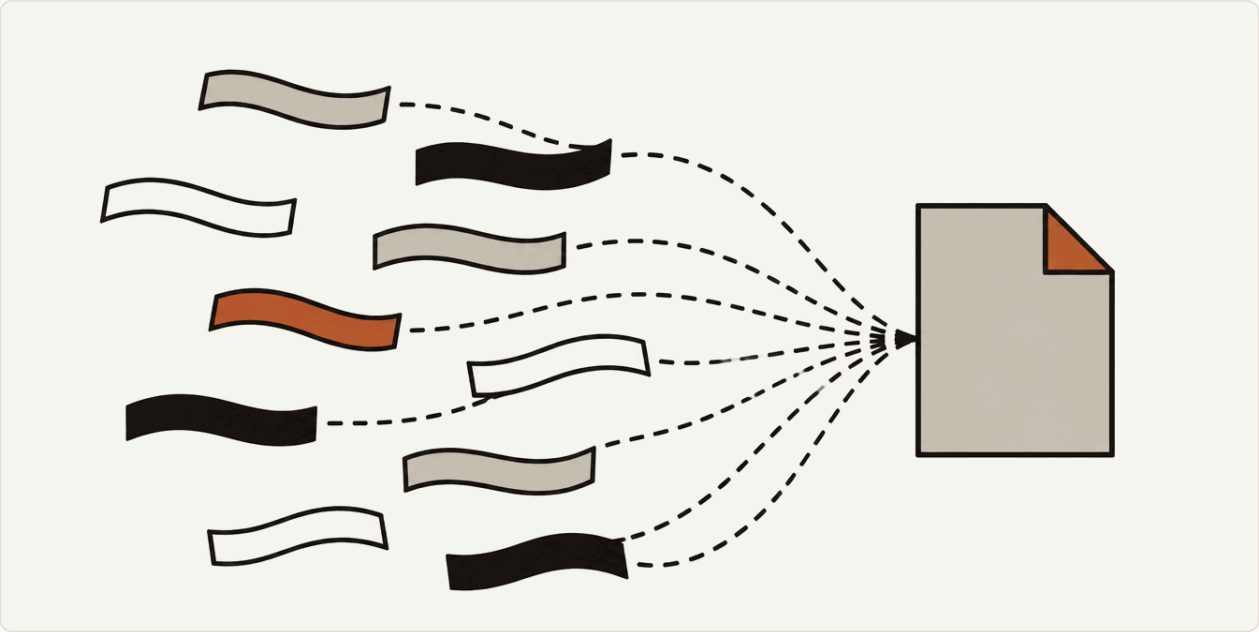
Tools change. AI products update often. If a screen looks different from the one described, the method still applies; owners get free updated versions in their library.

© 2026 Neurova AI. All rights reserved.

CHAPTER 01

What a research agent is, and isn’t

Staying current means reading the same sites every week, mostly to find out that nothing changed. The agent you'll build reads them for you all week, keeps a log of everything it saw, and puts one page in your inbox on Monday morning. Every line on that page links to where it came from.



YOU'LL HAVE	A clear design for your research agent, and a decision on whether to build or buy
TIME	20 minutes

Think of the sites you check every week to keep up: the statistics office, a few news searches, your trade association, maybe a central bank. Most weeks, most of them have nothing new for you. You still have to look to know that.

AI can do the looking. It can also get things wrong while sounding sure. In October 2025, journalists from 22 public broadcasters in 18 countries checked more than 3,000 answers about the news from ChatGPT, Copilot, Gemini and Perplexity. **45% had at least one significant issue**, and the most common one was sourcing: the answer didn't say where a claim came from, or said it wrongly (EBU and BBC, *News Integrity in AI Assistants*). That study tested assistants answering questions from their own searches, not summaries of articles you hand them, which is what you'll build. But the lesson carries over: keep the source behind every line, and keep a record of what was read.

Gather, judge, log, brief

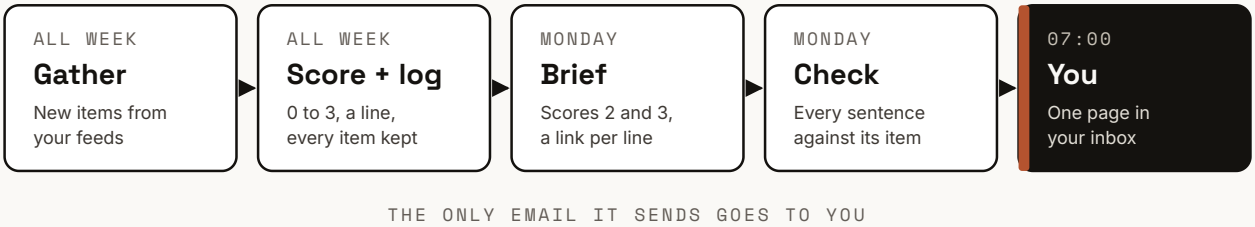


Fig 1.1 Items flow in all week; on Monday only what matters comes out, each line with its link

Your agent does four jobs, built as two automations:

1. **Gather:** pick up every new item from your sources as it's published.
2. **Judge:** score each item from 0 to 3 against a card that says what matters to you, with a one-line summary.
3. **Log:** write every item into a sheet, the dropped ones too, with its score and the reason.
4. **Brief:** on Monday morning, turn last week's best items into one page, check it against the items, and email it to you.

> THE THREE RULES
>
> **A link behind every line.** If the briefing says it, you can click to where it came from. **It informs, you decide.** The briefing reports what happened; it doesn't tell you what to do. **It logs what it read.** When you wonder why something was left out, the log tells you.

Build it, or buy it?

OPTION	GOOD AT	WHERE IT FALLS SHORT FOR THIS JOB
Google Alerts	Free emails when new pages match your search words	No judgement and no summary: you still read every alert
Scheduled tasks in ChatGPT and similar assistants	A daily or weekly briefing on your interests, written for you	You don't choose the sources or see what it read; it searches the web its own way
Feed readers with AI summaries	Many sources in one place, summarised	Another subscription, and another place to go and read
A paid market report or newsletter	Expert selection and comment	Their topics, not yours, on their schedule
Build your own (this guide)	Your sources, your scoring rules, a log, one email on your schedule	About five hours to build, and a paid Zapier plan

> WHEN BUYING IS THE BETTER CHOICE
>
> If a trade association or analyst already sends a weekly update that covers what you need, read that. Build your own when your mix of sources is your own, when you want to see what was left out and why, or when the sources are in a language your team doesn't read every day.

PROMPT 1.1 · IS A WEEKLY BRIEFING WORTH IT FOR ME? COPY · PASTE · FILL THE MARKED PARTS

Here are the sources I check to keep up with my market, and how often: `list, one per line`.
I spend about `minutes` minutes a week on them, and I mostly look for `what you look for`.
Estimate how much of that time a one-page Monday briefing would save, and which of my sources an automated agent
probably can't read (logins, paywalls, PDFs, social media). Then tell me honestly whether building one is worth it.

WHAT NEXT

Get the full guide

You've read the first chapter. The full guide takes you through all 6 chapters of the build, step by step, with every prompt ready to copy and a checklist to work from.

02 **Set up in an hour**

A card that says what matters, a tested source list, a reading log and a copyright check, before any automation.

03 **Build 1: the collector**

One automation that picks up every new item from your feeds, drops obvious noise for free, scores the rest against your card and writes every one into your log.

04 **Build 2: the Monday briefing**

A second automation that wakes up on Monday morning, reads last week's rows from your log and one page without a feed, writes one page, checks it against the items and emails it to you.

05 **Test it, and trust it**

Score real items you labelled before the first run, check one real briefing line by line, and plant errors to see whether the check catches them.

06 **Run it safely**

What it costs in tasks, the order to switch it on, five minutes on Friday with your log, a troubleshooting table, and what to check before you forward a briefing to anyone.

SEE THE FULL GUIDE

neurova-ai.com/guides/first-ai-agent/weekly-report/