

Hermes Agent – Prompt Reference Sheet

CONTENTS

What's inside

01	Prompt reference Hermes Agent resources	p.03
02	Day one Initial Setup	p.03
03	Week one Calibrate your agent	p.05
04	Month One Refine your agent	p.06
05	Month Three Review your agent	p.07

HERMES AGENT RESOURCES

Prompt reference

Every prompt you need to set up, calibrate, and maintain your Hermes research agent — in the order you'll use them.



Work through these prompts in order as your agent matures. The placeholders in italics are the parts you swap out for your niche and client. Copy each prompt directly into your Hermes Agent chat.

INITIAL SETUP

Day one

Research agent setup

Send this once to configure your agent and start the learning loop

COPY

You are a niche research agent for [describe your client's industry or focus area].

Every day, scan the following sources and log a summary of the most relevant findings:

Sources:

- [Source 1 — e.g. a specific newsletter, RSS feed, or subreddit]
- [Source 2]
- [Source 3]

For each item you log, include:

- The source and headline
- A 2-3 sentence summary
- Why it might be relevant to [client's niche or goal]

Save everything to intelligence_log.md.

After each daily scan, ask me: "What was useful today, and what wasn't?" Use my answer to adjust what you prioritise in future scans.

What this does

Sets up the daily scan routine and activates the learning loop. That last line is the key — it tells Hermes to collect feedback and use it to improve over time.

What you need to add:

- ☐ **[industry or focus area]**
e.g. "B2B SaaS", "sustainable fashion", "seed-stage startups"
- ☐ **[Source 1–3]**
specific newsletters, subreddits, RSS feeds, or blogs in the niche
- ☐ **[client's niche or goal]**
the outcome your client cares about, e.g. "finding funding opportunities"

CALIBRATE YOUR AGENT

Week one

Daily feedback

Send this at the end of each day during the first week to train Hermes

COPY

Today's scan: the piece on [topic or headline] was useful because [reason]. The [source name] summary wasn't relevant because [reason] — note that.

What this does:

This short daily message is what trains Hermes during week one. Volume is normal at this stage — your job is just to tell it what matters and what doesn't. One message per day is enough.

What to remember:

☐ **[topic or headline]**
the specific piece that stood out from today's report

☐ **[reason]**
why it was or wasn't useful for your client

☐ **[source name]**
the source that isn't pulling its weight

— REFINE YOUR AGENT

Month One

Calibration Review

Run this before your next scan to surface patterns and refine Hermes's priorities

COPY

We've been running for about a month. Before your next scan, I want to do a quick calibration review. Go through the last 30 days of `intelligence_log.md` and tell me:

1. Which sources produced the most useful results? Which ones were consistently rated as low-quality, off-topic, or repetitive?
2. Which stories or topics did I engage with most? What does that tell you about what I actually care about?
3. What are you still uncertain about in terms of my priorities — what would help you filter better going forward?

Based on your answers, suggest 2-3 adjustments to how you're running the scans. I'll tell you which ones to apply.

What this does:

This gets Hermes to surface its own patterns before you or your client review them. Send the output to your client as a calibration report — their job is to confirm what's right and redirect where Hermes is off track. This monthly check-in also helps you catch source rot early, before degraded sources quietly lower the quality of the reports.

— REVIEW YOUR AGENT

Month Three

Intelligence review

Run this before your weekly scan to check signal quality and share progress with your client

COPY

You've been running for 3 months. Before this week's scans, review `intelligence_log.md` — which sources have been consistently high-signal? Which topics keep surfacing? Are there any patterns I should know about? Then run your usual scan.

****What this does:**** By this point, Hermes should understand the niche well enough to flag outliers on its own. Send the output to your client as a progress report so they can see how the agent's priorities have evolved, and confirm what to keep, drop, or adjust. After six months, you'll rarely need to intervene — the agent will know the niche, the client's preferences, and which sources to trust.



****Keeping costs under control:**** Worth doing at the three-month mark to stop token costs from creeping up. As the `intelligence_log.md` file grows, so does the token spend per scan. At the three-month mark, switch to a "summary of summaries" approach: instead of feeding the full history back into every prompt, ask Hermes to pull from the last 1-2 weeks of the log only. You can also check your running costs any time using the `/usage` command inside Hermes Agent.

REFLECT

Jot down what you're taking with you.

Use this space for the ideas, questions, and next steps you want to remember.

Ready to deploy your Hermes agent?

Get Hermes up and running in less than 5 minutes with Hostinger's one-click autodeployment. Select your VPS plan, add credits from nexos.ai, and your Docker container will be ready to go. Use code APP10 for a discount.

Get started →

https://www.hostinger.com/vps/docker/hermes-agent?utm_source=youtube&utm_medium=organic&utm_campaign=hostingeracademydownloadables&utm_content=