FREE STARTER GUIDE · PREVIEW OF THE HERMES AGENT CASEBOOK

AI Agent Foundations Checklist

Build your first useful AI agent the right way: clear owner, boundary, input, output and stop condition — before the first prompt.

For	Level	Time	Next level
Beginners and AI operators	Beginner	45–60 minutes	Chapter 2 of the full playbook

THE MICRO-PAIN THIS GUIDE FIXES

Everyone is building agents. Most of them do nothing useful.

You've seen the demos: agents that research, write, book and code on their own. So you connect a model to a few tools, give it a big goal, and watch it loop, hallucinate or stop halfway. It feels like the technology isn't ready — but usually the design isn't.

80%
Of failed first agents fail because the goal was too broad, not because the model was weak.

Loops
Agents without a stop condition burn time and tokens repeating steps.

Trust
One confident wrong action and nobody wants to use the agent again.

When you finish this free guide you will have

- One well-scoped agent job written on a single card
- Clear inputs, outputs, tools and boundaries
- A plan-first instruction so the agent thinks before it acts
- A verification step and a handoff note after every run

THIS FREE STARTER

- One working route for one problem
- 7 short steps and a copy-and-use prompt
- One worked example and a final checklist

THE FULL PLAYBOOK · $29

- 6 chapters that take the whole job end to end
- 10 ready-to-use templates, a worked case, 126 PDF pages
- **See The Hermes Agent Casebook →**

STEP BY STEP

The working method

Five foundations — owner, input, output, boundary, stop — plus a plan-first habit and a verification step.

You need: Any agent-capable tool: Claude (Projects, Cowork or Claude Code), ChatGPT agents, n8n, or similar; A notes doc for the agent card; Sample inputs for testing. **Rule:** One agent, one job. Split big goals into several small agents.

STEP 01 OF 07

Pick a job, not a dream

- Choose a task that is repetitive, rule-based and low-risk.
- Describe it in one sentence starting with a verb.
- Check: could a new assistant do it from a one-page instruction?

Example: "Summarise each new customer feedback form into category, sentiment and one quote."

Done when: The job fits in one sentence and has an obvious 'done'.

STEP 02 OF 07

Name the owner

- Write who receives the output and who fixes problems.
- Decide how often the owner reviews results (every run, daily, weekly).
- Make the owner the only one who can change the agent's instructions.

Example: Owner: Head of Support, reviews summaries every morning.

Done when: One named owner and a review rhythm.

STEP 03 OF 07

Define input and output precisely

- Input: where it comes from, what format, what fields matter.
- Output: exact format (table, JSON, email draft), fields, length, destination.
- Include one example of a perfect output in the instructions.

Example: Input: form text. Output: row with category (billing/bug/feature), sentiment (-1/0/1), quote ≤25 words.

Done when: A sample input and ideal output exist side by side.

STEP 04 OF 07

Set the boundary

- List allowed tools and data sources only.
- List forbidden actions: send, delete, pay, publish, share externally.
- Start with read-only access plus a draft output.

Example: May read the feedback sheet; may write to a 'summaries' tab; may not email anyone.

Done when: Allowed and forbidden lists written.

> **In the full playbook** · This guide previews Chapter 2 of The Hermes Agent Casebook: “Scope each use case as a contract with a trigger, inputs, output, stop point and approval gate.” The chapter gives you the full method, the templates and a worked case. See what is inside · Get it for $29

STEP 05 OF 07

Add a stop condition

- Stop when the output is produced and self-checked.
- Stop and ask when input is missing, ambiguous or outside scope.
- Set a limit on steps or time.

Example: "If the form is empty or not in English, mark 'needs human' and move on."

Done when: Done, blocked and limit conditions all written.

STEP 06 OF 07

Make it plan before it acts

- Instruct: restate the goal, list unknowns, propose steps, then act.
- For the first runs, approve the plan manually.
- Remove approval only after several correct plans.

Example: Agent: 'Goal: summarise 12 forms. Unknown: 2 forms have no text. Plan: skip those, flag them.'

Done when: The plan step appears in every run log.

STEP 07 OF 07

Verify and hand off

- Ask the agent to check its output against the format and example.
- End each run with a note: what was done, what was skipped, what needs a human.
- Owner reviews a sample and updates instructions when a pattern of errors appears.

Example: "12 processed, 2 flagged (empty). Category uncertain for #7."

Done when: Every run ends with a handoff note.

COPY AND USE

Prompts and template

Paste into your AI tool and replace everything in [brackets].

PROMPT 1 · AGENT CARD (SYSTEM INSTRUCTION)

```
Role: You are an agent that [job in one sentence].
Owner: [name]. Output goes to: [destination].
Input: [source + format].
Output: [exact format + fields]. Example:
[example output]
Allowed tools: [list]. Never: [send/delete/pay/publish/share].
Before acting: restate the goal, list unknowns, propose your steps.
Stop and flag 'needs human' if input is missing, ambiguous or out of scope.
After acting: check your output against the format and end with a handoff note (done / skipped /
needs human).
```

WORKED EXAMPLE

A feedback-summary agent in one afternoon

Jess spent 3 hours a week reading customer feedback forms and tagging them.

1. Job: summarise each new form into category, sentiment and one quote.
2. Owner: Jess, reviewing every morning.
3. Input: form responses sheet. Output: 'summaries' tab, fixed columns.
4. Boundary: read the sheet, write one tab, nothing else.

RESULT

3 hours a week saved, with a clear record of every flagged item and an agent the team trusts.

FINAL CHECKLIST

Before you call it done

- [ ] Job written in one sentence
- [ ] Owner and review rhythm named
- [ ] Input source and format defined
- [ ] Output format with a perfect example
- [ ] Allowed tools listed
- [ ] Forbidden actions listed
- [ ] Done / blocked / limit conditions set

In the full playbook · The Hermes Agent Casebook also shows you how to scope each use case as a contract with a trigger, limits and a metric, and ships the checklists and templates to do it. See what is inside · Get it for $29

WHEN IT GOES WRONG

SYMPTOM	LIKELY CAUSE	FIX
Agent loops or never finishes	No stop condition	Add done/blocked/limit rules
Outputs look different every time	Output format vague	Specify fields and include a perfect example
Agent did something unexpected	Too many tools connected	Reduce to an allow-list; start read-only
Nobody fixes errors	No owner	Name one owner with a review rhythm

LOCKED IN THE FULL PLAYBOOK · TOOLKIT PREVIEW

This free starter shows the method. The Hermes Agent Casebook ships the ready-to-use files:

- Agent pilot plan template
- Case extraction card
- Hermes FAQ
- Hermes installation and configuration guide
- Hermes prompt library
- Hermes setup checklist

Unlock all templates — $29 one-time →

WHERE THIS FREE GUIDE STOPS

You fixed one problem. The full system fixes the job around it.

This starter gave you one route through one step. It deliberately leaves out what makes the result hold up at work:

- **The steps before and after this one.** The Hermes Agent Casebook covers the whole job in six chapters, not one slice of it.
- **Ready-to-use templates.** 10 copy-and-use files, so you do not rebuild the structure every time.
- **A full worked case** that shows every decision on a realistic project, start to finish.
- **Checks for when it fails**, so a result you ship tomorrow is still right next month.

YOUR NEXT STEP · THE HERMES AGENT CASEBOOK

Put a persistent AI agent to work on recurring tasks, safely, and prove it pays off.

What you will be able to do:

- Choose persistent-agent use cases by the recurring work they remove
- Scope each use case as a contract with a trigger, limits and a metric
- Install and configure Hermes isolated and least-privileged, in 15 phases
- Run a two-week shadow-to-live pilot, then decide to scale, fix or kill

6	10	1	126
chapters	templates	worked case	PDF pages

You just read a preview of **Chapter 2: Scope each use case as a contract with a trigger, inputs, output, stop point and approval gate**.

~~$49~~ **$29** one-time · instant access in your CosynkAI Library

Get the playbook — $29 one-time → ***See what is inside***

WORKING ON SEVERAL AI WORKFLOWS?

Complete Guide Library — all 8 CosynkAI playbooks for **$97** one-time ($472 combined regular prices). Already own a playbook? Its price is credited toward the Library.

Get the Library — $97 → Compare all 8 playbooks

You received this free guide after sending the keyword **AGENT**. Developed by CosynkAI.com · Edition 1.0 · September 2026. Educational resource, not an implementation service. Tools, prices and limits change; check current documentation before relying on a tool in production. Questions: hello@cosynkai.com · Privacy Policy · Digital Product Terms