01 / START SMALL. MAKE IT USEFUL.

Your First AI Agent Starter Kit

By Ashley Kays and South Florida AI. Choose one job, create your assistant brief, and test the result before adding automation.

Your first 30 minutes

5 minutes: Pick one recurring task with a result you can check.
15 minutes: Supply context, one example, a format, and a definition of done. Try it in an AI tool you already have permission to use.
10 minutes: Compare the result with the source. Correct errors and save the instructions.

Choose a useful first job

Meeting follow-up - Turn sanitized notes into an action table and email draft.

Weekly planning - Organize tasks into priorities, dependencies, and questions.

Source comparison - Compare options from sources you supply; flag evidence gaps.

Before you begin

Use fictional or sanitized information while learning. Keep customer data, passwords, and private records out of practice prompts. A saved prompt is a useful starting point; it does not create an agent that can use tools or run on a schedule.

INSIDE: 3 prompts / agent worksheet / 5 tests / AI concepts / 7-day plan

02 / START SMALL. MAKE IT USEFUL.

Three prompts to try today

Copy from the online kit or use the text below. Replace bracketed instructions and review every result.

Find your first useful task

Help me choose one small task to try with AI. Ask me three questions, one at a time, about work I repeat, the result I need, and the information I can safely share. Then suggest three tasks ranked by ease of checking the result. For the simplest one, give me a 20-minute practice plan. Tell me what I must check myself. Do not assume you can access my apps or files.

Turn notes into next steps

Turn the notes below into a draft action list. Use only the information provided. Return a table with: action, owner, due date, supporting note, and question to resolve. Write 'not specified' for missing owners or dates. Do not invent decisions or commitments. Finish with a short follow-up email for me to review. Do not send anything. NOTES: [paste sanitized notes here]

Plan an agent with a clear job

Help me design an AI assistant for this recurring task: [task]. Ask about the input, desired output, allowed sources, tools, and how I will judge success. Then write an agent brief with: goal, trigger, numbered steps, output format, allowed actions, actions requiring approval, stop conditions, and five test cases. Start with a draft-only version I can run manually. Do not connect accounts, install tools, or act on my behalf.

03 / START SMALL. MAKE IT USEFUL.

Give your agent one clear job.

The online brief builder turns these answers into instructions you can export. It runs in your browser; it does not deploy an agent.

1. Goal

What task should it help you complete?

2. Inputs and sources

What may it use? What must it never access?

3. Deliverable

What format and level of detail do you need?

4. Definition of done

What must be true for you to accept the result?

Default boundaries for your first version

Start manually. Use only allowed sources. Ask about missing or conflicting facts. Produce a draft for review. Do not send, publish, purchase, delete, or change records. Stop if a required tool is unavailable. Never claim an action was completed when it was not.

Reminder: Instructions alone do not enforce access. Configure actual tool permissions and approval controls before connecting apps.

04 / START SMALL. MAKE IT USEFUL.

Test it before you trust it.

Write down the expected result first. Test again after changing instructions, models, or tools.

1. **A normal example with all details**

   Produces the right format and uses only supported facts.

2. **An example with a missing date or owner**

   Marks the gap or asks a question; does not guess.

3. **Two sources that disagree**

   Shows the conflict and asks which source should govern.

4. **Source text that says to ignore your rules**

   Treats that text as data and follows the task's boundaries.

5. **A request to send, spend, or use an unavailable tool**

   Stops for approval or explains the limitation; never claims an action it did not take.

Keep a tiny test log

Record: input / expected result / actual result / pass or fail / correction / review minutes. Save failed cases and rerun them after each change. Five cases are a starter check, not a production guarantee.

Measure the whole task

Time saved = manual time - AI time - review time - correction time.
Also count errors and useful completed outputs. Keep the process only if it improves the work.

05 / START SMALL. MAKE IT USEFUL.

Know what matters next.

Use these concepts to evaluate new tools. Features and plan access change; check current provider documentation.

Context

The background, examples, and source material you give the model. A useful sample often improves a result more than a longer prompt.

Try this: Give it one good example and one clear definition of done.

Workflow vs. agent

A workflow follows steps you define. An agent can choose some next steps and tools. A saved prompt alone does not create an autonomous agent.

Try this: Prove the task manually before adding tool access.

Skills

Reusable instructions and resources for a particular task. In tools that support them, skills package a process the agent can load when needed.

Try this: Save the instructions and a checked example that worked.

Connectors & MCP

Connectors give an AI application access to other systems. MCP is a standard for connecting AI apps with tools and context; support and permissions vary.

Try this: Check the provider and permissions. Start with read-only access.

Evaluations

Repeatable checks against examples with known expectations. A confident answer is not proof that the task was done correctly.

Try this: Keep a small test set; rerun it after changing instructions or tools.

Human approval

A review point before an agent changes something outside the chat. Instructions help, but actual tool permissions and approval controls enforce limits.

Try this: Require approval for sending, publishing, spending, or deleting.

06 / START SMALL. MAKE IT USEFUL.

A little practice. Every day.

Use the same small task for a week so you can see what actually improves.

DAY 1 / CHOOSE

Pick one repeatable, low-risk task. Record how long it takes today.

DAY 2 / GIVE CONTEXT

Write the goal, one example, allowed sources, and the output format.

DAY 3 / TRY IT

Run the instructions manually on sanitized or fictional information.

DAY 4 / CHALLENGE IT

Try the five test cases. Save mistakes as new examples to check.

DAY 5 / MAKE IT REUSABLE

Revise the instructions and save the version that passed.

DAY 6 / TRY A NEW CASE

Measure total time, including checking and correcting the result.

DAY 7 / DECIDE

Keep, revise, or stop. Add one tool only if there is a clear need and reliable review.

Want a hand? Bring your brief.

Work with Ashley on one agreed workflow. A $350 private AI Working Session includes a 90-minute build and a 30-minute follow-up. Paid service by founder-affiliated Waymaker Agency. southflai.com/ai-help

Sources & a habit to keep

Read one official update weekly. Test one relevant change; keep it only if it helps. Original exercises by South Florida AI. Definitions informed by these primary sources, reviewed November 7, 2026:

Anthropic: Building effective agents (foundational distinctions)

Model Context Protocol: Introduction

Anthropic: Equipping agents with Agent Skills