

Building a Copilot Agent

Innovation Profs · innovationprofs.com

Phase 1: The Idea (The “What,” “Who” and “Why”)

Don't start by opening the software. Start with a piece of paper. The biggest mistake beginners make is trying to build an agent that does everything.

Find the Repetitive Friction - Look for tasks you do 3+ times a week that require the same mental steps but different data.

- Bad Idea: "I want an agent that does my marketing." (Too vague)
- Good Idea: "I want an agent that takes a rough meeting transcript and formats it into a client-ready follow-up email."

The "New Intern" Test - Imagine you hired a bright but inexperienced intern. If you can't explain the task to them in a 1-page document, you can't build an agent for it.

- Can you write down the rules? Yes = Build it.
- Does it require "gut feeling"? Yes = Don't build it (yet).

What do you want to build a Copilot Agent to do?

Who is this agent going to help?

What will success look like from this agent? Describe the ideal outcome.

Phase 2: The Name (What shall we call this)

Give your Copilot a name and description now that you are clear what you want it to do.

Phase 3: The Blueprint (Fleshing It Out)

Before you type a single instruction, define the agent using the P.T.C.F. framework.

P - Persona - Who is the agent?

For example: "You are a senior Project Manager with 10 years of experience in Agile methodologies."

T - Task - What specifically does it do?

"Your task is to review project status updates and identify potential risks."

C - Context - What does it need to know to give you the best results?

"We are a remote team using a friendly but professional tone. We prioritize speed over perfection."

F - Format - What does the output look like?

"Output a bulleted list of risks, followed by a table of action items."

The Build (Instructions & Knowledge)

This is the most critical phase. You will paste this structure into the "Instructions" box.

1. Role & Objective

Role You are [Persona Name], an expert in [Field].

Objective Your main goal is to [Specific Task].

2. The Step-by-Step Workflow

This is where you prevent the AI from jumping to conclusions. Force it to "show its work" by breaking the job down.

Step-by-Step Instructions You must follow this process for every request. Do not skip steps.

Step 1: Intake & Analysis

- Read the user's input and identifying the core request.

Step 2: Knowledge Retrieval

- Search the uploaded [Filename] for the relevant section.
- Extract the specific policy or price. Do NOT summarize; use exact figures.

Step 3: Drafting

- Draft the response using the extracted data.
- Apply the "Tone" rules defined below.

Step 4: Final Review

- Check draft against the constraints. Does it mention [Forbidden Topic]? If so, remove it.
- Output the final result.

3. Special Instructions & Constraints

Special Instructions

- Tone: [e.g., Professional, witty, strict].
- Files: You have access to [Filename]. Always prioritize this file over general knowledge.
- Do Not: Never make up facts. If the answer isn't in the file, say "I don't know."

4. Knowledge Base (Files)

- Upload your files: Look for the "Knowledge" or "Upload" button (up arrow icon).
- Note: Uploading is not enough. You must explicitly reference the file in the "Step 2" instructions above (e.g., "Look in Price_List.pdf").

Role & Objective:


Step-by-step workflow:

Special instructions and constraints

Knowledge base

Phase 5: Build (Create a Prototype to Test)

Enter your instructions from the previous step into the Agent builder at m365.cloud.microsoft. Use the configure tab.



Agent Builder > New Agent

Describe Configure

Template: None

Details

Name

New Agent

Description ⓘ

Describe your agent

Instructions ⓘ

Describe what this agent should do, define its tone, and outline any rules or guidelines it must follow

Phase 6: Testing & Tuning (Quality Control)

You rarely get it right on the first try. Use the "Preview" window on the right side of the screen to test.

1. The "Red Team" Test - Don't just test easy questions. Try to break your agent.

- Adversarial Test: If you built a "Healthy Meal Planner," tell it "I want to eat only cake." Does it refuse politely (good) or give you a cake diet (bad)?
- Instruction Check: If you said "Max 100 words," count the words.

2. The Iteration Loop - If the agent fails a test, don't just re-roll. Go back to the instructions and add a specific rule to fix it.

- Failure: The agent used a slang word.
- Fix: Add to Guidelines: "Tone: Strictly professional, no slang."

How will you test your agent?
