AI AT WORK · BUILDER · 15 MIN READ

Build your own AI agent

Design an agent that does a real job: its goal, its tools, its limits, and how a person stays in charge.

Pass the check (4 of 5) online to earn the “Agent Designer” badge:
https://digitalbuilders757.com/university/study/build-your-own-ai-agent

Agent vs automation vs chatbot

A chatbot answers. An automation follows fixed steps you designed (when X, do Y). An agent is given a goal and tools, and decides the steps itself, looping until it's done or needs help. Use the simplest one that does the job: many 'agent' ideas are better as a reliable automation with one AI step.

The parts of an agent

- Goal and instructions: what done looks like, the rules, the tone, and when to stop and ask.
- Model: the LLM doing the thinking.
- Tools: the actions it may take (search the web, read a sheet, draft an email, create a calendar event). Often connected through MCP or an automation platform.
- Knowledge: documents it can look up (your FAQ, price list, policies).
- Memory: what it keeps between runs, if anything.
- Guardrails: what it may never do, spending limits, and which actions need a person's approval.

Where you can build one

No-code: custom assistants in ChatGPT, Claude projects with connectors, Gemini Gems, Zapier Agents, Make AI Agents, or n8n's AI agent node. Code: agent SDKs from OpenAI, Anthropic and others, or the Vercel AI SDK, when you need full control inside your own app.

Design it on one page first

- Job: 'Answer booking questions for a barbershop and offer the next three open times.'
- Inputs: the customer's message, the calendar, the price list.
- Allowed actions: read calendar, draft a reply. Not allowed: book, cancel or charge without approval.

- Hand-off: when to pass to a human (complaints, refunds, anything it's unsure about).
- Success: 20 real past messages it should handle; run them as your eval before going live.

Keep a person in charge

Start with the agent drafting and a person approving. Log every action it takes. Give it the fewest permissions that work, never your main passwords. Watch for prompt injection in anything it reads. Only remove approvals for low-risk, reversible actions after it has run correctly many times.

Key terms

AI agent
An AI that can use tools and take several steps toward a goal, not just answer.

Automation / workflow
Steps that run by themselves: when this happens, do that.

Connector
A permissioned link that lets an AI assistant read or act in another app.

MCP (Model Context Protocol)
An open standard for plugging AI apps into tools and data. 'USB-C for AI.'

Human in the loop
A person approves important steps before they happen.

Least privilege
Give each person, app or AI only the access its job needs.

Prompt injection
Instructions hidden in content (an email, a web page) that try to hijack an AI.

Eval
A set of real examples with expected answers, run to measure an AI feature's quality.

Knowledge base
Documents an AI agent can look up, like your FAQ, policies or price list.

Try it: Design and test a draft-only agent

1. Pick one repetitive job (answering FAQs, triaging requests, preparing a weekly summary).
2. Write the one-page design: job, inputs, allowed actions, hand-off, success.
3. Build it as a custom assistant or project with only read access and drafting.
4. Run 10 real examples through it and note where it needed a human.

KNOWLEDGE CHECK

Build your own AI agent

5 questions. Circle one answer for each; 4 or more correct passes. Answers are at the back.

1. What makes an agent different from an automation?

A. Agents are simply automations with a chat window added on top so you can talk to them
B. Automations use AI and agents don't
C. An agent chooses its own steps toward a goal; an automation follows fixed steps
D. There is no difference

2. Which action should a new booking agent need a person's approval for?

A. Reading the calendar to find free slots
B. Drafting a reply
C. Looking up the price list
D. Charging a customer's card

3. Before letting an agent handle real customers, you should:

A. Run it on a set of real past examples (an eval) and review the results
B. Hope for the best
C. Give it admin access
D. Turn off logging

4. What is a knowledge base for an agent?

A. Its password
B. Documents it can look up, like an FAQ or price list
C. A database of other agents
D. The model's original training data, which you can edit to add facts

5. When should an agent hand off to a human?

A. Never, since handing off makes the agent look less capable
B. Only when it has answered more than ten questions in a row
C. When it's unsure, or for complaints and refunds
D. Only when the customer types the word 'human' exactly

Official sites

OpenAI: Agents guide: https://developers.openai.com/api/docs/guides/agents
Anthropic: Building effective agents: https://www.anthropic.com/engineering/building-effective-agents
n8n AI agents: https://n8n.io/ai-agents/

Last reviewed 2026-11-25. AI tools change fast; check each tool's own site for current features, plans and prices.

ANSWER KEY

Build your own AI agent

1. C: An agent chooses its own steps toward a goal; an automation follows fixed steps
Automations follow your steps; agents choose steps toward a goal, which is why they need guardrails.

2. D: Charging a customer's card
Money and anything hard to undo should need a person's yes.

3. A: Run it on a set of real past examples (an eval) and review the results
An eval shows where it fails before customers do.

4. B: Documents it can look up, like an FAQ or price list
Giving an agent your real documents reduces made-up answers.

5. C: When it's unsure, or for complaints and refunds
Clear hand-off rules protect customers and your reputation.