How to n8n Tutorial for Beginners: Build Your First Free AI Age

AI Automation & Agents · about 24 minutes · 2026-11-11

A true beginner's n8n tutorial: install it free, learn the canvas with a tiny practice workflow, connect a free Gemini model, and build your first AI agent with memory and tools in 30 gentle steps.

Things you will need

- A computer with 8 GB of RAM and 2 GB of free disk space
- Docker Desktop on Windows or macOS, or Docker Engine on Linux
- n8n community edition, self-hosted and free
- A Google account and a free Gemini API key from AI Studio
- Optional: a free Groq or OpenRouter account as a backup provider
- A Telegram or Gmail account for delivering the agent's output
- A password manager for the owner account and the API keys
- Sixty to ninety minutes for the first complete build

Understand n8n and set up your workspace

1. **Understand what n8n and an AI agent actually are**
   n8n is a free, fair-code automation tool: you drag boxes called nodes onto a canvas and connect them with lines, and data flows from one to the next.
2. **Choose how to run n8n: cloud, Docker, or npx**
   You have three free ways to start.
3. **Install Docker Desktop the easy way**
   Download Docker Desktop from the official site and run the installer with the default options; on Windows accept the prompt to enable WSL 2 when it appears.
4. **Launch n8n with a persistent volume**
   Open a terminal and create a storage volume with docker volume create n8n_data, then start n8n with docker run -it --rm --name n8n -p 5678:5678 -v n8n_data:/home/node/.n8n docker.n8n.io/n8nio/n8n.
5. **Create your owner account at localhost:5678**
   When the terminal settles, open http://localhost:5678 in your browser and create the owner account n8n asks for on first launch.
6. **Tour the editor: canvas, nodes, and panels**
   Spend two quiet minutes looking at the empty workflow before building anything.
7. **Set your timezone before you schedule anything**
   Stop the container and restart it with -e GENERIC_TIMEZONE=Europe/Berlin -e TZ=Europe/Berlin added to the docker run command, swapping in your own zone.

Build a tiny workflow to learn the canvas

8. **Start with a Manual Trigger**
   Before any AI, build a two-node practice workflow so the canvas becomes familiar.

9. Add an Edit Fields (Set) node

Click the plus to the right of the trigger, search for Edit Fields, and add it; older tutorials call this the Set node.

10. Run it and read the output JSON

Press the test button and watch both nodes turn green, then click the Edit Fields node to see its output: a small block of JSON containing your greeting field.

11. Understand items, fields, and expressions

Here is the one concept that unlocks everything else.

12. Name, save, and duplicate your workflow

Click the workflow title at the top and rename it to My first practice, then press save.

Connect a free AI model

13. Create a free Gemini API key

Sign in at Google AI Studio with any Google account and create an API key; the Gemini API free tier needs no credit card and is generous enough for a personal agent.

14. Save the key as an n8n credential

In n8n, open Credentials in the left menu, choose Add credential, search for Google Gemini(PaLM) Api, and paste the key there rather than into any node field.

15. Add the AI Agent node

Start a fresh workflow, click the plus button, search for AI Agent, and add it; it lives in the Advanced AI section and is the brain of this build.

16. Attach the Gemini Chat Model

Click the Chat Model slot under the agent and choose Google Gemini Chat Model, then select the credential you saved and a Flash-family model such as gemini-flash-latest.

17. Write your first system prompt

Open the agent's Options, add a System Message, and write the job in plain English: the role, the task, the output format, and the limits.

18. Test the agent in the chat panel

Open the agent node and use its built-in chat panel to type a simple question such as Explain what a workflow trigger is in one sentence, then run it.

Give your agent memory and tools

19. Add Simple Memory so it remembers

Click the Memory slot under the agent and add Simple Memory, then set the context window to around ten messages.

20. Add a Wikipedia tool for facts

Click the Tool slot under the agent and add the Wikipedia tool, which lets the model look up factual articles instead of guessing from memory.

21. Add a Calculator tool for numbers

Add the Calculator tool from the same Tool slot, because language models guess arithmetic and dates unless a tool computes them.

22. Add an HTTP Request tool for web pages

Add the HTTP Request Tool so the agent can fetch the full text of a URL you provide, which turns it into a simple summarizer of any page you paste.

23. Write tool descriptions the model can read

The most overlooked setting is each tool's description, because the model chooses tools by reading descriptions, not by inspecting their configuration.

24. Set safety limits: iteration cap and read-only tools

Open the agent's Options and set maximum iterations to five or ten so a confused model cannot loop forever and burn your free quota in minutes.

Trigger, deliver, and keep it running

25. Add a Chat Trigger to talk to your agent

Add a Chat Trigger node and connect it to the agent's input, then click Open Chat to talk to your build in a proper chat window instead of the node panel.

26. Deliver answers to Telegram or Gmail

Add a Telegram node after the agent to push answers to your phone, or a Gmail node to email them to yourself; Telegram is simpler because a bot token from BotFather is all it needs.

27. Add a Schedule Trigger for a daily run

When you want the agent to work without you, add a Schedule Trigger set to every day at 08:00 as an alternative to the chat trigger.

28. Read execution logs to debug failures

Open the Executions tab, click the latest run, and inspect each node's input and output by clicking it on the canvas.

29. Activate and verify one unattended run

Flip the Active toggle in the top-right corner, then wait for the next scheduled time and confirm a successful execution appears without you touching anything.

30. Back up, secure, and stay inside free limits

Finish with three habits that keep the build alive.

Sources

1. Install n8n with Docker
   https://docs.n8n.io/hosting/installation/docker/
2. AI Agent node documentation
   https://docs.n8n.io/integrations/builtin/cluster-nodes/root-nodes/n8n-nodes-langchain.agent/
3. Advanced AI in n8n
   https://docs.n8n.io/advanced-ai/
4. Gemini API key documentation
   https://ai.google.dev/gemini-api/docs/api-key
5. Gemini API pricing and free tier
   https://ai.google.dev/gemini-api/docs/pricing
6. Docker Desktop overview
   https://docs.docker.com/desktop/
7. n8n community forum
   https://community.n8n.io/

Written as a checklist from https://www.howcompass.com/how-to/n8n-Tutorial-for-Beginners-Build-Your-First-Free-AI-Agent. Corrections: https:/