

# Local AI Smoke Test

*Install a local model. Run three prompts. Verify nothing left your machine.*

---

There's a category of your life and work that should never touch a cloud server — and the only way to handle it is AI that runs entirely on your own hardware. This week you install Ollama or LM Studio, pull one model that fits your machine, and run three real prompts with the network watching. You verify, with your own eyes, that nothing left your machine.

## THE ACTION

- Install Ollama or LM Studio (the Sovereign Stack page walks the options).
- Pull one model sized to your hardware:
  - 32GB+ unified RAM: Qwen 3.6 or Gemma 4 26B-class.
  - Below that: Phi-4 Mini or a small Llama.
- Run three real watchman prompts entirely locally.
- Verify nothing left the machine — turn off wifi mid-session and confirm it still answers.

### SOVEREIGNTY IS PROVEN, NOT ASSUMED

"It runs locally" is a claim until you've watched it answer with the network off. Prove it once with your own eyes, and then you can trust the boundary.

## SCRIPTURE

*Proverbs 22:3 — "A prudent man foresees evil and hides himself." Some things you hide on your own hardware before there's ever a breach to react to. That's prudence, not paranoia.*