

LEARN AI · FREE GUIDE 02

Build Your First AI Agent

The anatomy of an agent, five starter builds ranked by difficulty, and the mistakes that kill v1 — from someone who has shipped agents into live operations.

What actually makes it an 'agent'

- LLM brain — reads the goal, decides the next step
- Tools — email, sheets, WhatsApp, ERP APIs, web search: the hands
- Memory — what it did, what worked, where it left off
- Loop — observe → decide → act → check, until done or blocked
- Guardrails — approval steps, spend limits, 'never do' lists

5 starter agents, easiest first

- Inbox triage agent — labels, drafts replies, escalates the ugly 10%
- Daily brief agent — pulls sheets/dashboards, writes your 7 AM summary
- Follow-up agent — chases pending vendor/candidate replies politely
- RFQ agent — sends structured quote requests, tabulates responses
- Report agent — turns raw exports into a formatted weekly review deck

Industry example — A 40-person ops team shipped 8 production agents after a 6-week cohort — the RFQ agent alone cut a 6-day procurement cycle to 1.5 days.

Your first build, step by step

- Pick ONE task you do weekly that follows rules you can write down
- Write the SOP as if for a smart intern — that text becomes the system prompt
- Start with human-in-the-loop: agent drafts, you approve, it sends
- Log every run for 2 weeks; fix the top 3 failure patterns only
- Remove the approval step only for the cases it has never got wrong

Why v1 agents die

- Too many tools on day one — start with two, maximum
- No definition of 'done' — the loop wanders and burns tokens
- Testing on clean data — feed it your real, messy Tuesday inbox
- Zero fallback — every agent needs a 'stuck? ping the human' branch
- Automating a broken process — fix the process first, then automate

Want this implemented, not just downloaded?

Book a free 30-min discovery call → consult@hariadvisory.com · +91 98801 27460