

NOT EVERY PROBLEM NEEDS AN AI AGENT


Sometimes the smartest solution is a rule, a workflow, a chatbot—or no AI at all.

SIMPLE TASK



Press switch → light turns on

OVERENGINEERED




The goal is not “more AI.” The goal is the right solution.

Do Not Hire a Detective to Flip a Switch


A simple, predictable job does not need planning, investigation, or autonomy.

RIGHT TOOL



**One action. One known
result.**

TOO MUCH




Plan → search → reason → act
~~→ check~~
~~More cost. More delay. More ways~~
~~to fail.~~

Complexity should solve a real problem—not create one.

Choose the Smallest Solution That Works

Move up only when the simpler level cannot handle the task.





LOWER AUTONOMY

HIGHER AUTONOMY

Add autonomy only when the task truly needs more flexibility.

Use a Rule When the Answer Is Exact

Rules are best when the logic is stable and the result must be predictable.




Do not ask an agent to guess what simple code can guarantee.

Use a Workflow When the Steps Are Known

A workflow follows a designed path. It can branch, but the branches are already defined.

EXAMPLE: LEAVE REQUEST




The same process happens every time. Exceptions can be sent to a person.

No free-form planning is required.

Use a Chatbot When You Need One Good Answer

Not every AI task needs tools, memory, or a loop.



EXPLAIN

"Explain photosynthesis simply."



REWRITE

"Make this email clearer."



SUMMARIZE

"Give me the key points."



BRAINSTORM


"Suggest 10 names."

Ask → answer → done. That is useful AI, even without an agent.

Use an Agent When the Path Can Change

The goal is clear, but the next step depends on what the agent discovers.

GOAL: RESTORE AN EMPLOYEE'S VPN




Same goal. Different next action based on evidence.

Quick Test #1: Can You Write Every Step in Advance?

This question separates many workflows from agents.

YES



Use a workflow

Example: Generate the monthly payroll file, then send it for approval.

NO



An agent may help

Example: Investigate why one customer cannot log in after an update.

“No” does not automatically mean agent. Check the next questions too.

Quick Test #2: Does the Task Need All Three?

Agents earn their place when several difficult capabilities must work together.




Judgment + Tools + Feedback = A strong reason to consider an agent

If the task needs only one of these, a simpler system may be enough.

The Hidden Cost of an AI Agent

More autonomy brings more capability—and more engineering work.

WORKFLOW



- Usually follows one known path
- Predictable speed and cost
- Easier to test and audit
- Fewer moving parts

AGENT



- May call models and tools many times
- Path and runtime can vary
- More failure points to observe
- Needs stronger evals and guardrails

More capable does not automatically mean better.

High-Stakes Decisions Need Human Control

For decisions that can seriously affect a person, AI should assist—not silently decide.



Require qualified human review for high-impact outcomes such as:

HEALTHCARE

Diagnosis or treatment

FINANCE

Loan or insurance decision

LEGAL

Rights or legal outcome

PEOPLE

Hiring, firing, discipline

Autonomy must match the risk.

Same Domain. Different Solution.

Choose based on the task—not the industry label.

SIMPLE SYSTEM

BOUNDED AGENT

HEALTHCARE

Appointment reminder

Summarize a complex case for
clinician review

BANKING

EMI calculator

Investigate an unusual fraud alert

EDUCATION

Send exam result email

Create a study plan that adapts
weekly

RETAIL

Order-status lookup

Resolve an unusual delivery
problem

A domain can contain rules, workflows, chatbots, and agents.

Full Example: Customer Support

Three requests arrive. Only one truly needs an agent.

A

“Where is my order?”




DATABASE LOOKUP

Read order ID → return status

B

“Reset my password.”



WORKFLOW

Verify identity → send reset link

C

“The app fails only on office Wi-Fi after yesterday’s update.”



BOUNDED AGENT

Investigate context →
search docs → test hypotheses → escalate

Use autonomy only where uncertainty requires it.

Before You Build, Check Need AND Readiness

An agent needs a strong reason to exist—and a safe way to operate.

DO WE NEED AN AGENT?

- ✓ The path cannot be fully predefined.
- ✓ The task needs multiple tools or systems.
- ✓ The next action must change after feedback.
- ✓ The value justifies extra cost and latency.

ARE WE READY?

- ✓ Success and failure are clearly defined.
- ✓ Permissions, limits, and approvals are set.
- ✓ We can trace, test, and evaluate behavior.
- ✓ There is a safe fallback or human escalation.

If readiness is weak, do not deploy more autonomy yet.

DAY 5 TAKEAWAY


USE THE LEAST AUTONOMY THAT SOLVES THE PROBLEM



EXACT ANSWER



RULE



KNOWN PATH



WORKFLOW



ONE RESPONSE



CHATBOT



CHANGING PATH



AGENT

Good engineering is not about using the most advanced architecture.

It is about choosing the simplest reliable system for the job.

NEXT UP

Day 6: How Agents Use Tools — Function Calling

Comment a task: rule, workflow, chatbot, or agent?