


THE AGENT LOOP


How an AI Agent Gets Things Done

Give it a goal. It decides, acts, checks the result, and adjusts.



The Easiest Analogy: A GPS That Reroutes

A GPS does not memorize one fixed route. It keeps checking what happens next.



- 1 Destination**
Reach the airport
- 2 Choose route**
Pick the best next path
- 3 Observe**
Notice traffic or a road closure
- 4 Adapt**
Reroute and continue

An agent works the same way—one decision at a time.

The Agent Loop, Step by Step

1  **GOAL**
What result are we trying to achieve?



2  **DECIDE**
What is the best next step?



3  **ACT**
Use a tool or produce an action.



4  **OBSERVE**
Read what actually happened.



5  **ADAPT OR STOP**
Continue, change course, ask, or finish.

Repeat until the goal succeeds, a limit is reached, or human help is needed.

Step 1: Give the Agent a Clear Goal

The loop needs a destination. A vague goal makes the agent wander.

TOO VAGUE



“Help me plan a trip.”

Where? When? Budget?
What counts as “done”?

CLEAR GOAL



**“Plan a 3-day Jaipur trip
under ₹25,000, leaving
Friday, with one heritage
hotel.”**

RESULT

What should exist?

LIMITS

Budget, time, rules

SUCCESS

How do we know it
worked?

PERMISSION

What may it do?

Step 2: Decide the Next Useful Action

A good agent does not blindly do everything at once.



Before acting, it checks:

- What do I already know?
- What information is missing?
- What is the smallest useful next step?
- Which action is safe and allowed?

TRIP EXAMPLE

- 1 Check dates and budget
- ↓
- 2 Search travel options
- ↓
- 3 Compare price and time
- ↓
- 4 Build the itinerary

Agentic does not mean random. It means choosing the next step from the current situation.

Step 3: Act

Using the Right Tool

The model decides. A tool lets the agent do something in the real world.



SEARCH

Find current information



CALCULATOR

Compare prices or totals



CALENDAR

Check or schedule time



FILES / DATA

Read and update records



CODE

Run tests or analyze data




MESSAGES

Draft or send—with permission

Tools turn a decision into an action.

Step 4: Observe What Actually Happened

The agent must read the result—not assume that the action worked.



ACTION

Check hotel availability

RESULT

"No rooms available."



SUCCESS

Booking confirmed



CHANGED

Price increased



FAILED

Tool returned an error



MISSING

More information needed

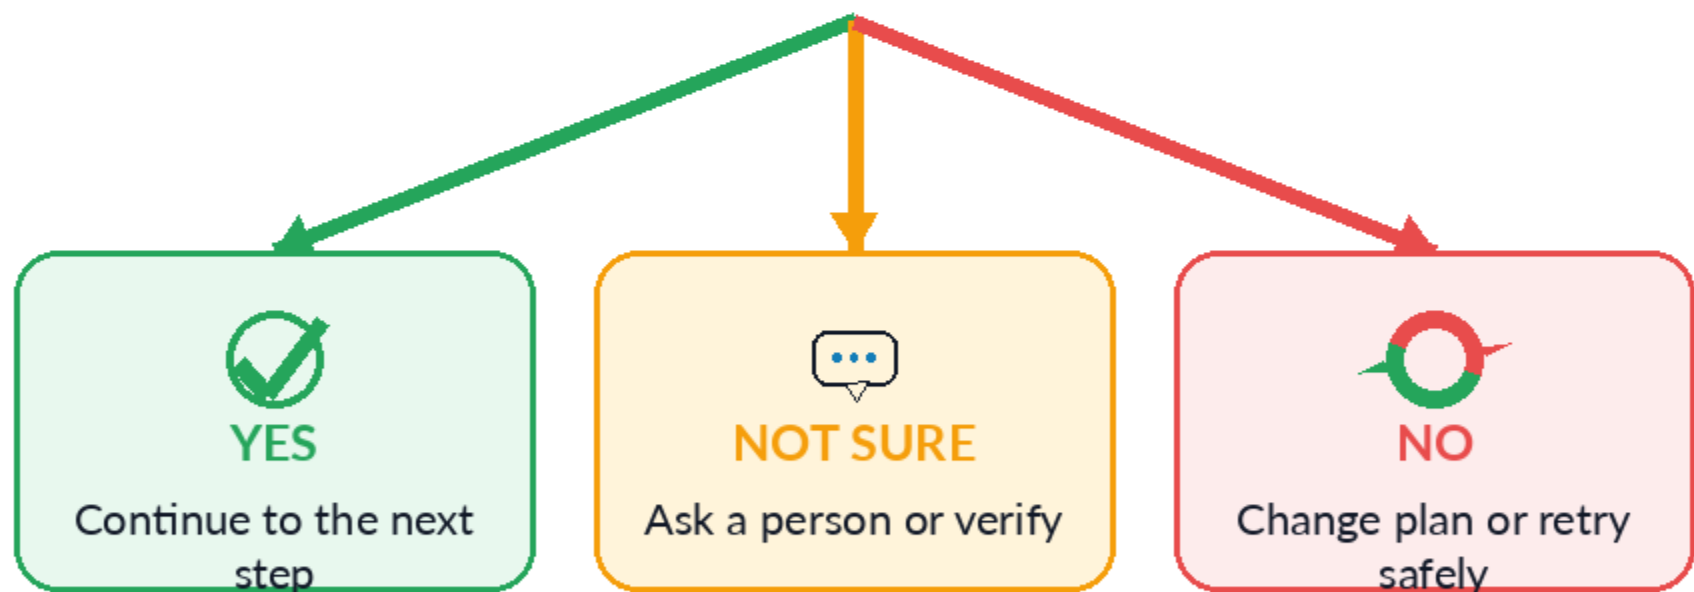
Observation is feedback for the next decision.

Step 5: Adapt, Ask, or Stop

After every result, the agent chooses what should happen next.

OBSERVATION

Did the last action move us closer to the goal?



Adaptation—not one-time tool use—is what makes the loop useful.

Full Example: Book a Dinner Table



GOAL: Book dinner for 4 tonight at 8 PM, in Jubilee Hills, under ₹4,000, with outdoor seating.

1



SEARCH

Find nearby
restaurants

2



FILTER

Rating, budget,
seating

3



CHECK

Availability for 8
PM

4



SHORTLIST

Keep the best
options

**Current best match: "Garden Table" – ₹3,600 estimated,
outdoor seating, 4.5/5**

Next action: check whether a table is actually available.

Same Goal. The Path Changes.



OBSERVATION:

"Garden Table is fully booked at 8 PM."

1

Do not pretend it succeeded



2

Check the next-best restaurant



3

Verify outdoor seating and price



4

Ask the user before booking



5

Book only after approval



The goal stayed fixed. The next action changed.

The Same Loop Works Across Domains

Only the goal, tools, and safety rules change.



CUSTOMER SUPPORT

Read issue → search knowledge
→ suggest fix → check whether
it worked



STUDENT TUTOR

Assess answer → choose lesson
→ ask a question → adjust
difficulty



SOFTWARE ENGINEERING

Inspect bug → run tests →
change code → rerun tests




OFFICE ADMIN

Check calendars → find options
→ ask approval → schedule
meeting

One pattern. Many professions.

An Agent Also Needs Rules for When to Stop

Without a finish rule, an agent may repeat actions forever.



NO STOP RULE

Search → fail → search
→ fail → search...



Result: wasted time, money,
and tool calls.

STOP WHEN...



Goal achieved



No progress after a safe retry
limit



Time or budget limit reached



Required permission is missing

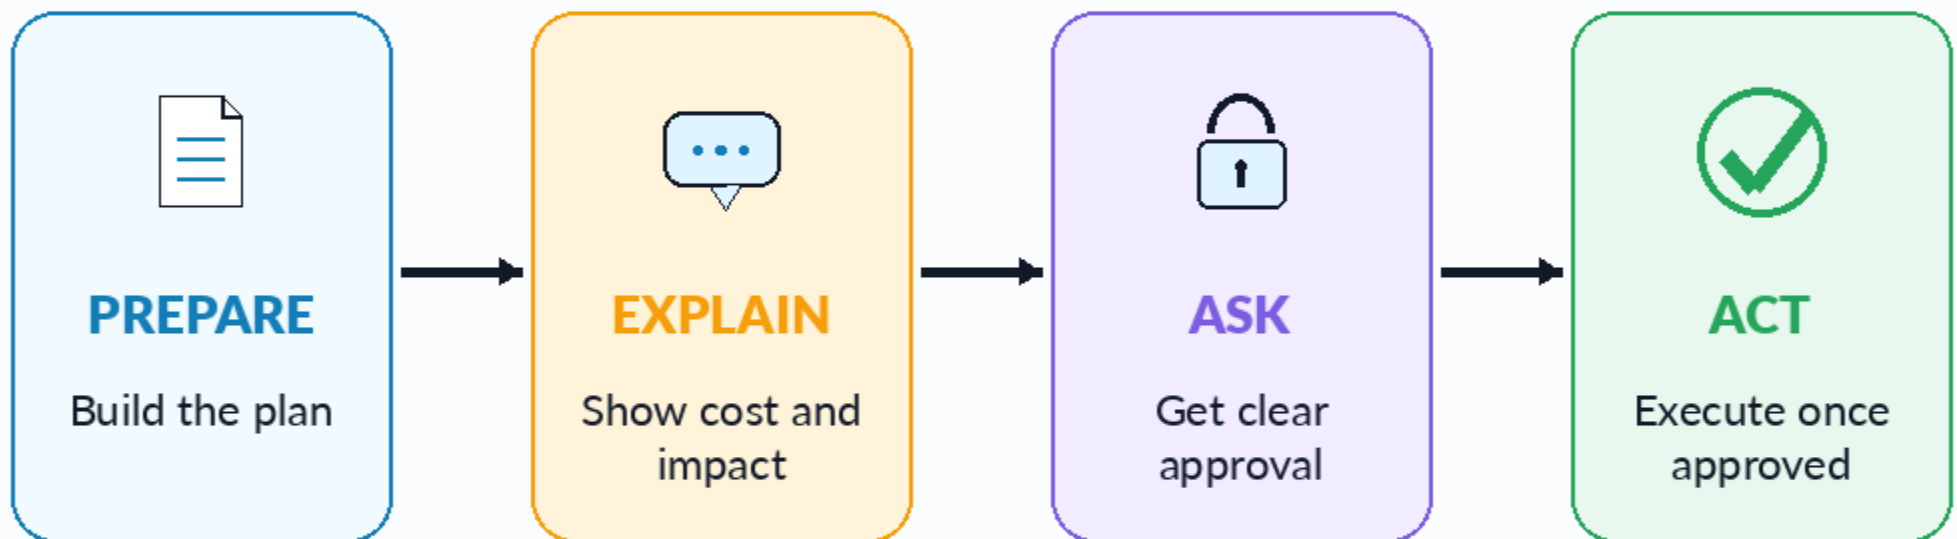


Confidence is too low—ask a
human

A good agent knows when to finish—and when to ask for help.

Some Actions Need Human Approval

An agent can prepare a risky action without being allowed to execute it.



Always pause before high-impact actions:



PAYMENTS



SEND
MESSAGES



DELETE
DATA



PRODUCTION
CHANGES

Autonomy should grow with trust—not replace accountability.

What Can Go Wrong Inside the Loop?

Most agent failures are not magic. They are engineering problems we can control.

FAILURE



Keeps repeating the same action



Reads a tool result incorrectly



Chases an impossible goal



Spends too much time or money



Attempts a risky action

PROTECTION



Maximum steps and retry limits



Validate structured outputs



Clear constraints and stop rules



Budgets, timeouts, and model routing




Permissions and approval gates

Reliable agents are built with limits, checks, and visibility.

DAY 4 TAKEAWAY


An Agent Is a Feedback Loop



A useful agent does not act once. It checks every result and chooses what to do next.



Clear goal + useful tools + feedback + safe stopping rules



NEXT UP

Day 5: When an AI Agent Is the Wrong Solution