SHOULD AN AGENT PLAN BEFORE IT ACTS?

Meet Agent Planning — deciding what to do, in what order, and when to change course.

GOAL: PLAN MY CONFERENCE TRIP

1. Find flights
2. Choose hotel
3. Check schedule
4. Ask approval
5. Book

Planning turns one goal into a sequence of decisions and actions.

NOT EVERY TASK NEEDS A PLAN

Planning has overhead. Simple tasks are often better handled directly.

ACT DIRECTLY

“What is 15% of ₹8,000?”

One clear step.
Low uncertainty.
No tool chain needed.

PLAN FIRST

“Research, compare and book a 5-day trip.”

Many dependencies.
Multiple tools.
Consequential actions.

Use the simplest control strategy that can reliably finish the task.

WHAT IS AGENT PLANNING?

A plan is an explicit or implicit representation of the steps needed to move from goal to completion.

GOAL → SUBGOALS → ACTIONS

- Understand constraints
- Gather missing evidence
- Choose next action
- Check result
- Continue / re-plan / stop

A plan is not the answer. It is the control structure for reaching the answer.

THE CORE LOOP: PLAN → ACT → OBSERVE → ADAPT

Useful plans are not static checklists.

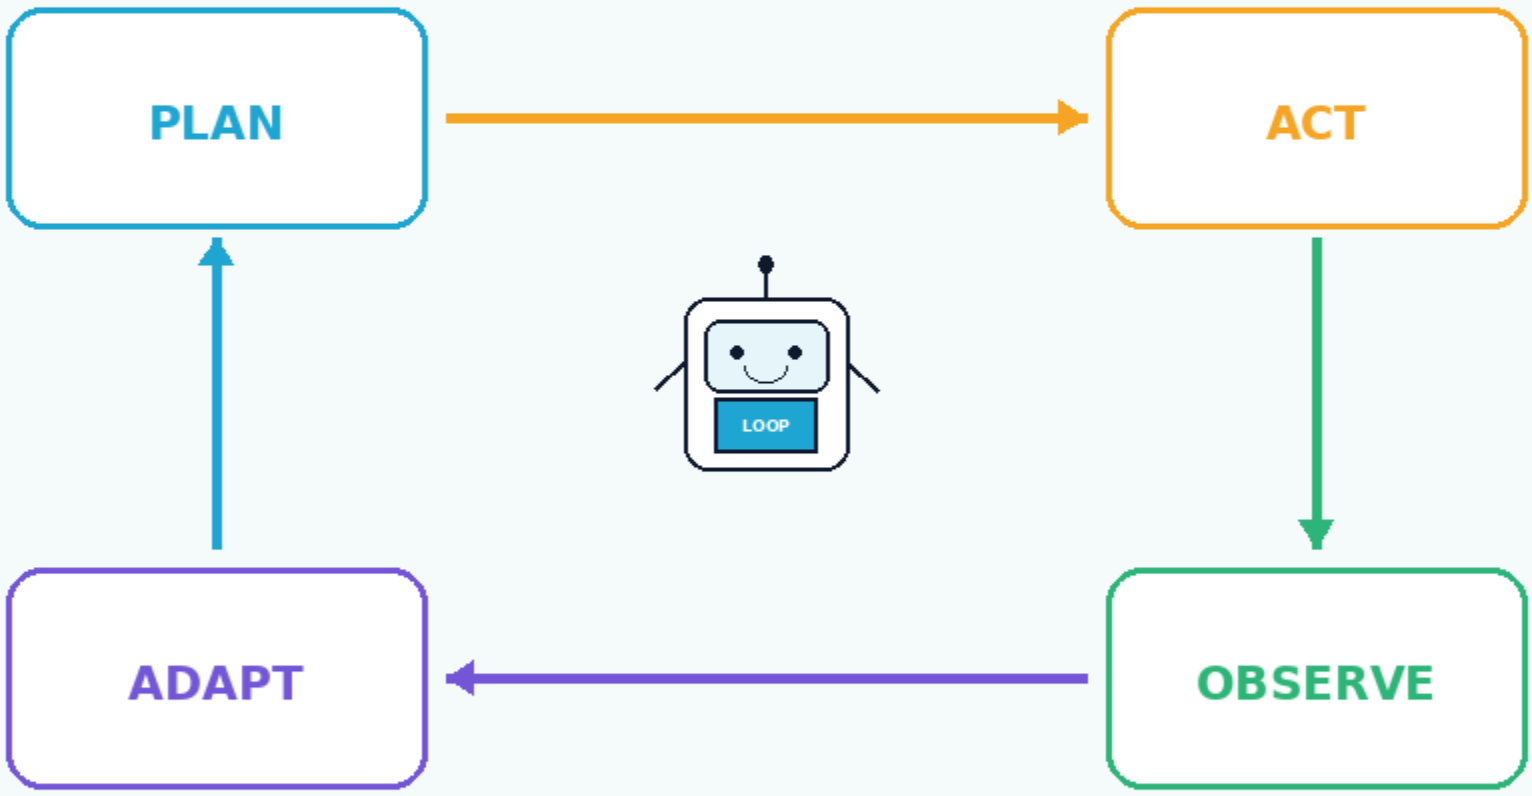

Execution produces new evidence. Good agents let that evidence change the plan.

STATIC PLAN VS DYNAMIC PLAN

The real world changes while the agent is working.

STATIC PLAN

Step 1
Step 2
Step 3
Step 4

Assumes nothing changes.

DYNAMIC PLAN

Plan
↓
Act
↓
Observe
↓
Re-plan if needed

A useful agent plan is usually a hypothesis — not a contract.

WHEN SHOULD THE AGENT PLAN?

Planning is most useful when the task has structure the model must discover or manage.

MULTI-STEP

DEPENDENCIES

MULTIPLE TOOLS

UNCERTAINTY

HIGH CONSEQUENCE

LONG-RUNNING

The more dependencies and uncertainty, the more valuable explicit planning becomes.

TASK DECOMPOSITION: BREAK THE GOAL INTO WORKABLE PARTS

Planning often starts by finding useful subgoals.

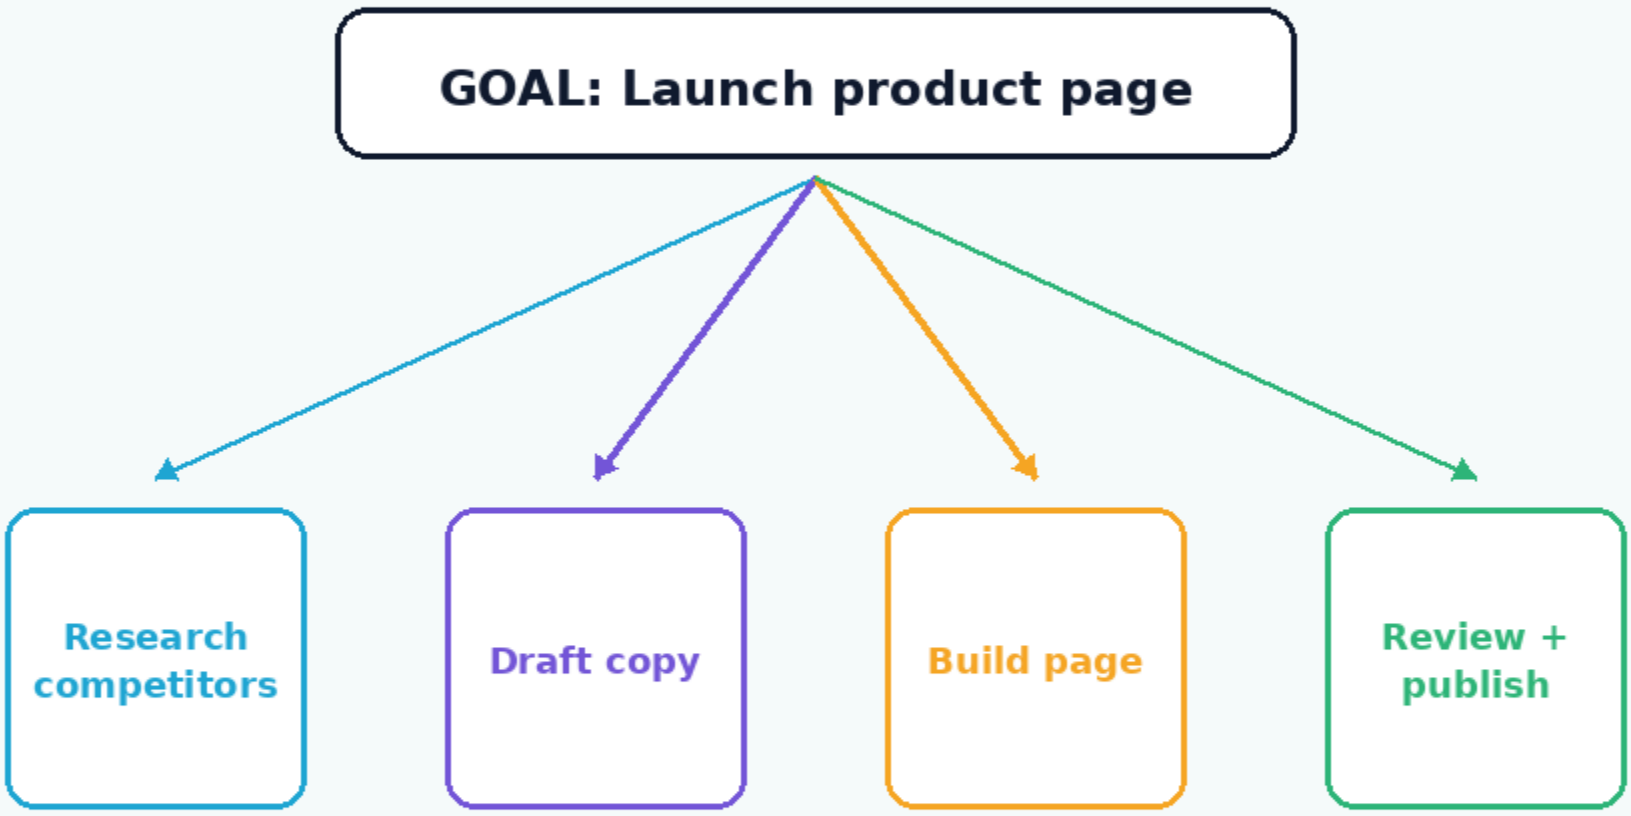

Good decomposition creates steps that are meaningful, checkable and actionable.

DEPENDENCIES DECIDE THE ORDER

Some steps can run in parallel. Others must wait.

SEQUENTIAL

PARALLEL

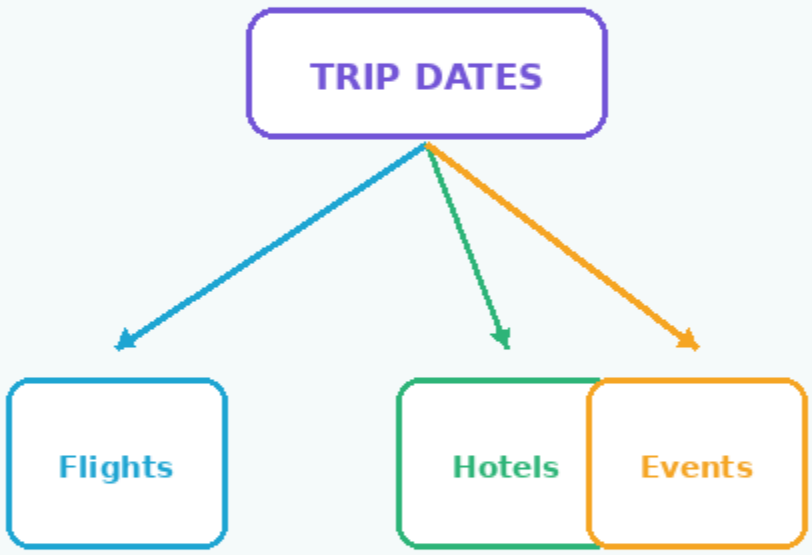

Planning is partly a dependency graph: what must happen before what?

PLANS NEED PRECONDITIONS AND SUCCESS CHECKS

A step should know when it can run — and how to know it worked.

BEFORE THE ACTION

PRECONDITIONS

✓ User dates known
✓ Budget available
✓ Search tool enabled

AFTER THE ACTION

SUCCESS CHECKS

✓ Results returned
✓ Constraints satisfied
✓ State updated

Without checks, a plan can keep moving even after a step silently failed.

RE-PLANNING: WHEN SHOULD THE AGENT CHANGE COURSE?

The next step should depend on what actually happened.

TOOL FAILED	→	RE-PLAN
NEW USER INPUT	→	RE-PLAN
CONSTRAINT CHANGED	→	RE-PLAN
NEW EVIDENCE	→	RE-PLAN
PLAN STUCK	→	RE-PLAN

Re-planning is a state transition triggered by evidence — not random indecision.

DO NOT OVER-PLAN

A 20-step plan for a 2-step problem is just latency, cost and brittleness.

OVER-PLANNED

1. Restate goal
2. Analyze goal
3. Analyze analysis
4. Build meta-plan
5. Re-check meta-plan

...

Too much ceremony.

LEAN PLAN

1. Get missing input
2. Execute tool
3. Verify result

Enough structure to be reliable.

Plan only to the level that improves decisions.

SOME PLANNING SHOULD LIVE IN CODE

Not every workflow decision belongs to the LLM.

LLM DECIDES

Open-ended choices
Ambiguous goals
Tool selection
Adaptive decomposition

CODE DECIDES

Hard business rules
Required approvals
Retry limits
Known workflow order

Mix model intelligence with deterministic orchestration when predictability matters.

PLANNING + TOOLS: THE PLAN MUST TOUCH REAL CAPABILITIES

A beautiful plan is useless if no executable action can realize it.

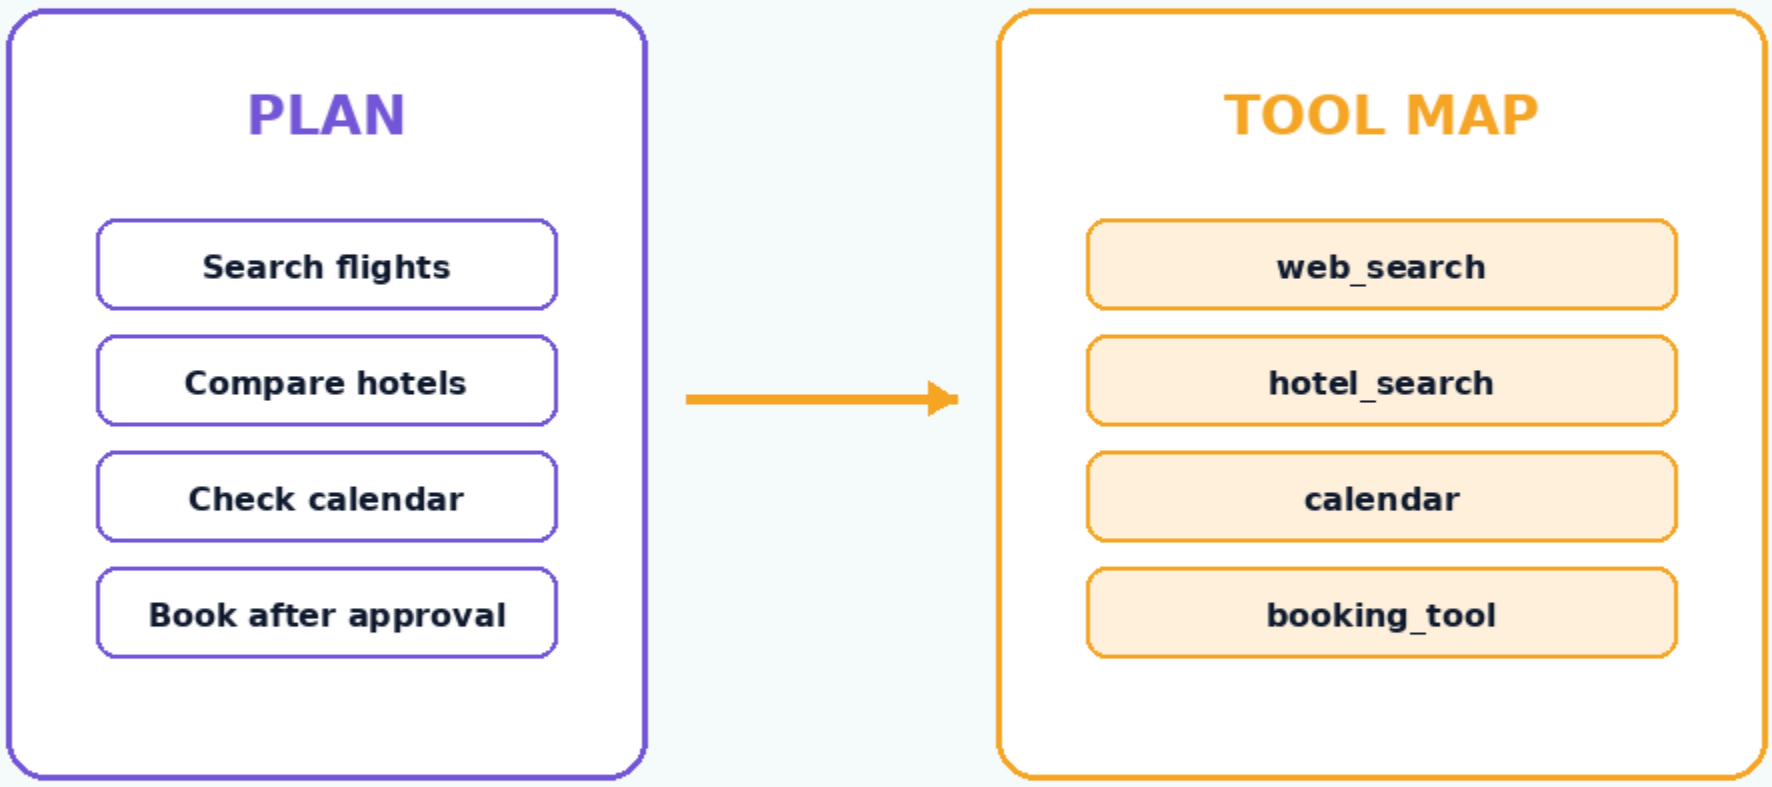

Every planned action should map to a real tool, agent, or deterministic operation.

PLANNING + MEMORY: USE WHAT YOU ALREADY KNOW

Memory can reduce planning effort by supplying constraints and prior lessons.

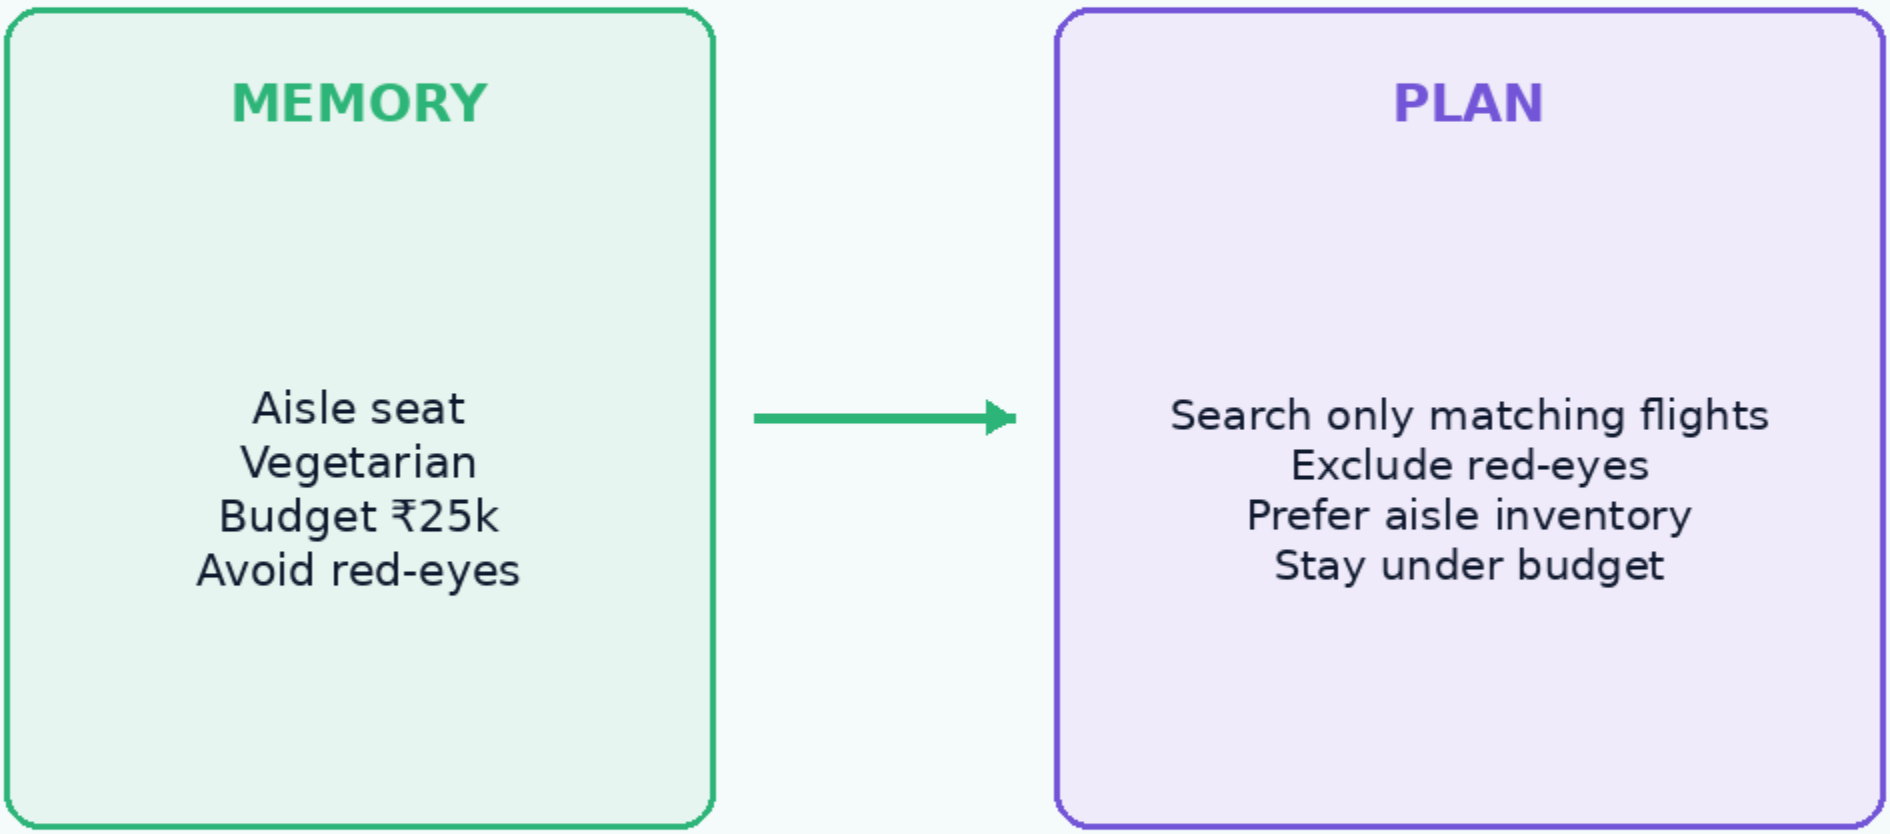

Better context produces smaller, more relevant plans.

PLANNING + EVALS: GRADE THE PROCESS, NOT ONLY THE ANSWER

Day 21 meets Day 23.

TASK SUCCESS

UNNECESSARY STEPS

TOOL CHOICE

RECOVERY

COST / LATENCY

POLICY COMPLIANCE

An agent can reach the right answer through a bad plan — until that bad plan fails.

REAL EXAMPLE: PLAN MY CONFERENCE TRIP

A good planner stays goal-directed while reacting to evidence.

1. Ask dates + budget
2. Search flights + hotels
3. Check conference schedule
4. Conflict found → re-plan
5. Present best itinerary
6. Ask before booking

The plan changes, but the goal and constraints remain stable.

A PRACTICAL PLANNING ARCHITECTURE

Planning is one layer inside the broader agent harness.

The planner proposes direction. The harness still owns execution boundaries.

DAY 23 TAKEAWAY

Agent planning is choosing and adapting a path from goal to completion.

GOAL | PLAN | ACT | OBSERVE | RE-PLAN | STOP

DECOMPOSE

Break goals into useful subgoals

ORDER

Respect dependencies

ADAPT

Re-plan when evidence changes

SIMPLIFY

Act directly when a plan adds no value

**THE BEST PLAN IS NOT THE LONGEST PLAN.
IT IS THE SMALLEST PLAN THAT KEEPS THE AGENT ON TRACK.**

NEXT: DAY 24 — Agent Reasoning Patterns: ReAct, reflection, critique and search.