

Ourbit Master AI - Day 8: Agents and Tool Use

Purpose: Understand plan > act > observe loops, tools, stop conditions and approval gates.

Source/build status: Recovered Module 5 is the source; preserve the weather example.

Learning outcomes

- Explain the plan > act > observe loop used in tool-enabled systems.
- Define allowed tools, permissions and stop conditions before consequential action.
- Handle tool failures without converting missing observations into fabricated facts.
- Match human approval requirements to the consequence of an action.

Core lesson - Agents and Tool Use

An AI agent can use a model in a loop: decide what should happen next, invoke an allowed tool or action, observe the result, then continue or stop. Plan > act > observe is useful shorthand, but real systems also need permissions, validation, stop conditions and error handling. Consequential actions require stronger authorization and human control.

Worked example

Goal: “Find today’s weather and summarize it.” Plan the lookup, call an approved weather tool, observe the result and decide whether enough information exists. If the tool fails, handle or surface the failure-do not fabricate data.

Teaching / video direction

Animate Plan > Act > Observe > Continue/Stop using the weather example. Add tool-failure and approval branches. Close by asking what changes when several actions and validations must work together-Day 9.

Hands-on practical exercise

- Choose a harmless task that requires one external tool or lookup.
- Define the goal, allowed tool and information the tool may access.
- Write the planned action and the observation required before continuing.
- Define success, failure and stop conditions.
- Identify any step that would require explicit human approval if the task became consequential.
- Describe what the system should do if the tool returns an error or no result.

Knowledge check

- What do plan, act and observe mean in an agent loop?

- Why are stop conditions necessary?
- How should a system respond when a tool fails?
- Why should tool permissions follow least privilege?
- Which kinds of actions deserve stronger human approval?
- How does tool use change the risk compared with text generation alone?

Assessment and remediation contract

The scored daily quiz uses 10 questions; 7/10 (70%) is a pass. Missed concepts map to exact lesson anchors for targeted remediation, including on a passing attempt. Attempts 1-9 below threshold return to remediation and a fresh eligible question set; an unsuccessful tenth attempt restarts the day in a new learning cycle while preserving history. The 50-question validated bank is a learner-release gate and will be scaled only after the Day 1 runtime passes acceptance.

Production / QA notes

- Use the worked scenario as the primary visual sequence; avoid generic AI-brain imagery.
- Keep terminology vendor-neutral and define unavoidable technical terms in context.
- Do not equate fluent output with verified truth or AI assistance with autonomous authority.
- Before learner release, fact-check the final lesson and validate every scored item against explicit teaching.

Learning comes before credentials. Completing this reading alone does not issue a certificate. Programme assessment and practical requirements apply.