

Ourbit Master AI - Day 14: Design an End-to-End AI Solution

Purpose: Combine prompting, context, retrieval, tools, evaluation, privacy and human gates into one defensible design.

Source/build status: Integration day.

Learning outcomes

- Translate a user problem into an end-to-end AI system with explicit success measures.
- Justify when prompting, retrieval, tools or automation are actually needed.
- Design evaluation, permissions, logging, escalation and human ownership into the architecture.
- Test the design against realistic failure and misuse cases before calling it viable.

Core lesson - Design an End-to-End AI Solution

An end-to-end AI solution is a controlled system, not a model surrounded by features. Begin with users, outcome and success measure; identify approved knowledge and data; choose prompting, retrieval or tools only where needed; define evaluation, permissions, logging, escalation and human ownership; then test the complete workflow against realistic failure cases.

Worked example

Design a vendor-neutral internal policy assistant: authorized users > current approved policy corpus > retrieval > grounded generation > visible source references > evaluation checks > access controls > escalation for insufficient or conflicting evidence > logging > accountable policy owner. Every component must solve a defined need.

Teaching / video direction

Build the architecture left-to-right, then run three failure cases through it: stale source, insufficient evidence and unauthorized request. Close by explaining that Day 15 is not another theory lesson-the learner must demonstrate an integrated solution and its evidence.

Hands-on practical exercise

- Choose a low-risk knowledge or workflow problem and define users, outcome and success measure.
- Draw the end-to-end flow from input through evidence/model/tool steps to final outcome.
- For every component, write why it exists and what happens if it fails.
- Add access controls, evaluation, logging, escalation and accountable ownership.
- Run at least three failure cases through the design, including insufficient evidence and unauthorized use.

- Revise the architecture based on the failures and record the changes.

Knowledge check

- Why is a model alone not an end-to-end solution?
- When is retrieval justified?
- Why must escalation behavior be designed before deployment?
- What should logs make possible?
- Why test unauthorized and insufficient-evidence cases?
- What evidence would show the solution improves the target task?

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.

Teaching / video script - 6-8 minute target

Open: Start with the worked scenario, not a generic definition: Design a vendor-neutral internal policy assistant.

Teach: Demonstrate the method visibly: Define users/outcome > approved knowledge > retrieval > generation > source display > evaluation > permissions > escalation > logging > human ownership. Pause at the judgment point and explain why the system should not simply continue on autopilot: Every component needs a reason. Include failure modes, fallback behavior and a way to test whether the system actually improves the target task.

Practice bridge: Ask the learner to reproduce the method on a different low-risk example, mark the verification/control step, and explain the reason for it.

Close: Day 15: demonstrate practical mastery in the capstone.

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.