

# **Ourbit Master AI - Day 7: RAG and Grounded Knowledge Workflows**

Purpose: Combine retrieval with generation while keeping answers traceable to supplied knowledge.

Source/build status: Recovered Module 4, second half, expanded for grounding.

## **Learning outcomes**

- Explain the retrieval-plus-generation pattern behind RAG.
- Keep generated answers traceable to selected source material.
- Recognize stale, irrelevant, incomplete or conflicting retrieval as failure modes.
- Respond appropriately when the available evidence cannot support an answer.

## **Core lesson - RAG and Grounded Knowledge Workflows**

Retrieval-augmented generation (RAG) combines retrieval with generation: retrieve relevant source material, supply selected material as context, then generate an answer. RAG is not an automatic truth machine; retrieval can select stale or irrelevant evidence and generation can go beyond it.

## **Worked example**

An employee asks whether unused annual leave carries forward. Retrieve the current approved policy and answer from the relevant clause. If evidence is ambiguous or conflicting, say so and escalate rather than inventing a rule.

## **Teaching / video direction**

Show Question > Retrieval > Selected passages > Generation > Answer with source reference, plus a branch for insufficient or conflicting evidence. Close by moving from answering with knowledge to bounded tool use-Day 8.

## **Hands-on practical exercise**

- Choose a short non-sensitive source document you are allowed to use.
- Write one question fully answered by it, one only partly answered and one not answered.
- For each question, identify the exact supporting passage or state that none exists.
- Draft answers that do not go beyond the available evidence.
- Add a source reference and an escalation or uncertainty statement where needed.
- Review whether any generated sentence introduces an unsupported claim.

## Knowledge check

- What two operations are combined in RAG?
- Why can a RAG answer still be wrong?
- What should happen when retrieved evidence does not answer the question?
- Why should source identity and version be retained?
- What is the difference between relevance and authority?
- How does Day 7 build on semantic retrieval from Day 6?

## 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.