

Ourbit Vendor-Neutral AI - Day 7: Retrieval-Augmented Generation

Day 7 - Retrieval-Augmented Generation

Programme: Vendor-Neutral AI - 15-Day Certification Programme

Purpose: Learn how retrieval can ground generation in selected knowledge while preserving evidence boundaries and uncertainty.

Learning outcomes

- Explain the retrieve > ground > generate > verify pattern.
- Distinguish model knowledge from retrieved evidence.
- Recognize missing, stale, irrelevant and conflicting retrieval.
- Design source traceability into answers.
- Specify safe behavior when evidence is insufficient.

1. Why retrieval is added

A language model can generate from its available context, but an organization may need answers based on a specific handbook, policy library or current knowledge base. Retrieval selects relevant material and supplies it as evidence for the generation step. This can improve grounding without retraining the model.

2. The four-stage pattern

Retrieve: find candidate evidence. Ground: make the relevant evidence available with its identity and boundaries. Generate: answer using that evidence and the task instruction. Verify: check whether important claims are actually supported. Each stage can fail independently.

3. Retrieval failure becomes generation risk

If retrieval misses the controlling policy, returns an obsolete document or selects a semantically related but irrelevant passage, a fluent answer can still be wrong. RAG does not turn weak evidence into truth. Retrieval quality, authority and freshness therefore belong in the system design.

4. Preserve traceability and uncertainty

Important claims should be traceable to the material used. If sources disagree, the system should expose the conflict or apply a documented authority rule. If evidence is missing, it should say that the available material is insufficient rather than filling the gap confidently.

Worked example - employee policy assistant

An employee asks about parental leave. The system retrieves an old handbook and a newer approved policy. A robust workflow uses effective date and approval status to select the controlling source, answers within that source's scope, and provides a reference. If jurisdiction or employee category changes the rule and that information is missing, the assistant asks for it or states the limitation.

Guided practice

Design a RAG flow for an internal handbook. Define the query, retrieval filters, evidence shown to the model, answer instruction, citation/traceability behavior, and what happens when no sufficiently authoritative source is found.

Knowledge check

- Why add retrieval to a generative workflow?
- What is the difference between retrieval and grounding?
- Name three retrieval failures.
- Why must claims remain traceable?
- What should happen when evidence is insufficient?

Applied evidence task

Two retrieved documents give conflicting answers to a customer-policy question. Produce a resolution design using authority, effective date, uncertainty handling and human escalation. Do not simply instruct the model to choose whichever document appears first.

Day 7 takeaway

RAG is an evidence workflow, not a truth switch. Strong implementations control what is retrieved, preserve source identity, verify claims and handle missing or conflicting evidence explicitly.

Assessment boundary

This learning evidence belongs only to vendor-neutral-ai-15-day.

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