

Enterprise RAG with Spring AI

Master Retrieval-Augmented Generation for your internal documents.

DURATION

1 day

LEVEL

Professional

GROUP RATE

\$2,900

CURRICULUM & HIGHLIGHTS

- RAG Architecture
- Vector Databases (PGVector)
- Embeddings & Vectorization
- Ingestion Pipelines
- Similarity Search

For an AI to be genuinely useful in a professional environment, it needs access to unique company knowledge. The challenge lies in retrieving the right data, vectorizing it, and presenting it to the model as relevant context at the right moment without compromising security or search performance today.

This workshop focuses entirely on the RAG architecture in Spring AI through Active Learning. We work practically with vector databases like PGVector, implement embeddings, and build document ingestion pipelines for your needs. Participants learn similarity search and prompt contextualization to deliver accurate, source-based answers based on the company's own documents and PDF files. We build a solution where your AI stops guessing and instead provides answers based on verified internal facts in a stable and scalable architecture for Enterprise environments here today.

TERMS

Fixed rate (max 12) for remote delivery. Extra costs for onsite.

YOUR OUTCOME

A functional RAG pipeline that can be searched and presented from your internal documents.

TARGET AUDIENCE

Teams wanting to use private data to build secure AI solutions in Spring Boot today.


PREREQUISITES

Basic knowledge of Spring AI equivalent to the Jumpstart workshop is required now.

RECOMMENDATION

Make your AI proactive with the 'AI Agents & Tools Foundations' workshop next.

TECHNOLOGIES & FOCUS



← INFO & INQUIRY



actlea.com

ACTIVE LEARNING EXCELLENCE

Andreas Olsson

EXPERT INSTRUCTOR | 20+ YEARS EXP