

Updated on 02/27/2026

Register

Weaviate training

2 days (14 hours)

Overview

Weaviate is an open-source vector database designed for semantic and hybrid search. It combines similarity search, structured filtering, and a native GraphQL API within an object-oriented architecture.

Our Weaviate training will help you understand how the platform works internally, configure its schema, and optimize hybrid search.

You will learn how to structure your data, adjust the weighting between lexical and vector search, and integrate Weaviate into a modern RAG pipeline.

By the end of this training, you will be able to design a high-performance, industrializable vector search architecture.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

Objectives

- Understand vector and hybrid search.
- Master the Weaviate object-oriented model.
- Optimize query relevance.
- Integrate Weaviate into a RAG system.
- Prepare for production deployment.

Target audience

- Data Engineers
- Backend developers
- AI engineers
- Data architects
- AI DevOps

Prerequisites

- Basic knowledge of Python and REST APIs
- Embedding concepts
- General knowledge of LLMs

Weaviate training

[Day 1 - Morning]

Fundamentals of vector and hybrid search

- Understanding embeddings and semantic search
- Vector search vs. lexical search
- Principles of hybrid search
- Similarity metrics: cosine, dot product, L2
- Weaviate's position in the AI ecosystem
- Hands-on workshop: Creating a class and inserting vectors.

[Day 1 - Afternoon]

Internal architecture and object-oriented model

- Object-oriented schema
- Classes, properties, and relationships
- Internal indexing
- Metadata management
- Native GraphQL API

Hybrid queries and optimization

- Top-k vector queries
- BM25 + vector combination
- Weighting and tuning
- Metadata filters
- Performance optimization
- Hands-on workshop: Optimizing a hybrid query.

[Day 2 - Morning]

Weaviate and RAG

- Architecture of a RAG pipeline
- Chunking and structuring data
- Integration with an LLM
- Native embedding modules
- Improving relevance and re-ranking
- Hands-on workshop: Setting up a mini hybrid RAG.

[Day 2 - Afternoon]

Advanced use cases

- Advanced document search
- Product recommendation
- Multimodal (text + image)
- Relevance management
- Versioning strategies

Industrialization and production

- Docker and cluster deployment
- Horizontal scalability
- Security and access management
- Observability
- Production best practices
- Hands-on workshop: Production-ready target architecture.

Target companies

This training is intended for both individuals and companies, large or small, wishing to train their teams in new advanced IT technologies or to acquire specific business knowledge or modern methods.

Placement at the start of training

The positioning at the start of the training complies with Qualiopi quality criteria. Upon final registration, the learner receives a self-assessment questionnaire that allows us to assess their estimated level of proficiency in different types of technologies, as well as their expectations and personal objectives for the upcoming training, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could be problematic for the monitoring and smooth running of the training session.

Teaching methods

Practical training: 60% practical, 40% theory. Training materials distributed in digital format

to all participants.

Organization

The course alternates between theoretical input from the trainer, supported by examples and discussion sessions, and group work.

Validation

At the end of the session, a multiple-choice questionnaire is used to verify that the skills have been correctly acquired.

Certification

A certificate will be issued to each trainee who has completed the entire training course.