

Updated on September 11, 2026

[Sign Up](#)

RAGFlow Training

3 days (21 hours)

Overview

RAGFlow is an open-source Retrieval-Augmented Generation platform that transforms complex documents into reliable, explainable AI assistants tailored to business needs.

Our RAGFlow training will help you master the essential steps in designing a RAG solution, from document ingestion to the generation of context-aware responses with their sources.

You'll learn how to configure knowledge bases, choose the appropriate content segmentation methods, select language and embedding models, and then monitor the quality of document indexing.

The training will also help you fine-tune hybrid search, similarity thresholds, reranking, and retrieval strategies to improve the relevance of responses while minimizing hallucinations.

Finally, you'll design conversational assistants and agents that leverage your domain expertise. You'll explore best practices for security, monitoring, maintenance, and deployment of a RAGFlow solution.

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

Objectives

- Understand the architecture and use cases of a RAG solution with RAGFlow
- Create and populate knowledge bases tailored to business documents

- Configure document segmentation, embeddings, and indexing to optimize search
- Adjust hybrid retrieval, reranking, and relevance settings
- Design a conversational assistant and an agent connected to document sources
- Prepare for the operation, security, and evolution of a RAGFlow platform

Target Audience

- AI Developer
- Data Engineer
- Solutions Architect
- MLOps Engineer
- AI Project Manager

Requirements

- Knowledge of LLM principles and generative AI
- Experience administering Linux or container environments
- Basic understanding of information retrieval, unstructured data, and APIs
- Basic experience with Python or an equivalent programming language

Technical Prerequisites

- Operational RAGFlow instance, deployed locally or in a shared environment
- Docker 24.0.0 or higher and Docker Compose 2.26.1 or higher for a standalone deployment
- An x86 machine with at least 4 cores, 16 GB of RAM, and 50 GB of disk space for self-hosting
- Configured access to an LLM provider and an embedding model

RAGFlow Training Agenda

[Day 1 - Morning]

RAG Fundamentals and Getting Started with RAGFlow

- RAGFlow's role in the generative AI ecosystem and business use cases
- Principles of Retrieval-Augmented Generation: ingestion, indexing, retrieval, and generation
- Platform architecture: services, document storage, search engine, and AI models
- Exploring the interface: workspaces, datasets, assistants, agents, and model settings for models
- Deployment with Docker Compose, service verification, and initial configuration of model providers

[Day 1 - Afternoon]

Building and preparing knowledge bases

- Creating a knowledge base and organizing corpora by business domain
- Importing structured and unstructured documents: PDFs, presentations, spreadsheets, office documents, and text content
- Selecting embedding models based on language, corpus, and quality constraints
- Selecting chunking methods: general, manual, book, table, presentation, and other document-based models
- Parsing verification, chunk editing, keyword addition, and management of active documents
- Hands-on workshop: creating a knowledge base, importing a domain-specific corpus, and validating the quality of document segmentation

[Day 2 - Morning]

Optimizing search and relevance

- How hybrid search works: full-text search, vector search, and result merging
- Adjusting the similarity threshold, vector weight, and number of results returned
- Using reranking to improve the order of sources and balancing latency and precision
- Indexing strategies: processed text, generated questions, enriched context, and file name weights
- Hierarchical segmentation, parent-child strategy, and contextual enrichment of retrieved

[Day 2 - Afternoon]

Retrieval tests and conversational assistant

- Designing test scenarios based on actual user questions
- Analysis of cited sources, extracted chunks, and causes of imprecise responses
- Iterative adjustment of parsing, keywords, search parameters, and reranking
- Creating a conversational assistant connected to one or more knowledge bases
- Drafting system instructions: business scope, response style, uncertainty handling, and citations
- Hands-on workshop: configuring an information assistant, measuring retrieval quality, and optimizing its responses using a set of business-specific questions

[Day 3 - Morning]

Agents and data pipelines

- Differences between RAG assistants, agents, and data-processing-oriented workflows

- Agent Design: Instructions, Memory, Knowledge Retrieval, and Output Control
- Exploring Pipeline Components: Parsing, Chunking, Enrichment, Indexing, and Routing
- Automating ingestion with data pipelines and synchronizing external sources
- Introduction to API and MCP integrations for connecting RAGFlow to business applications

[Day 3 - Afternoon]

Scaling, governance, and deployment

- Target architecture: sizing, choosing a document engine, persistence, and service configuration
- Securing access to models, documents, and sensitive data
- Knowledge lifecycle management: updates, deactivation, synchronization, and traceability
- Monitoring of assistants and agents: execution logs, error analysis, and Quality
- Deployment best practices: environment separation, backups, model evolution, and pre-production validation
- Hands-on workshop: finalizing an end-to-end RAGFlow solution, defining its and present validation criteria

Target Audience

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

Assessment at the Start of Training

The pre-training assessment complies with Qualiopi quality standards. Upon final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency with various types of technologies, as well as their expectations and personal goals 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 pose challenges for monitoring and ensuring the smooth running of the training session.

Teaching Methods

Hands-On Training: 60% Practical, 40% Theoretical. Training materials distributed in digital format to all participants.

Organization

The course alternates between theoretical input from the instructor—supported by examples and reflection sessions—and group work.

Assessment

At the end of the session, a multiple-choice quiz is administered to verify that participants have successfully acquired the skills.

Certification

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