



PostgreSQL AI Vector Search and RAG Applications

16 – 20 November 2026
London

PostgreSQL AI Vector Search and RAG Applications

Ref: 103600692_83904 **Date:** 16 – 20 November 2026 **Location:** London **Fees:** 5700 Euro

Course Overview:

PostgreSQL with AI: Vector Databases, RAG and Intelligent Applications is a practical PostgreSQL AI Training course for professionals building secure, data-grounded solutions. Participants progress from PostgreSQL fundamentals, administration, and AI-assisted SQL Training to PostgreSQL embeddings, pgvector Training, and vector similarity search. The course explains exact retrieval, HNSW and IVFFlat indexing, metadata filtering, and PostgreSQL Semantic Search before connecting these capabilities to language models. Participants then design Retrieval-Augmented Generation pipelines and complete PostgreSQL Chatbot Development exercises using governed enterprise data.

Grounded in PostgreSQL documentation, pgvector implementation guidance, RAG research, and recognized AI risk practices, this PostgreSQL AI Course balances development with PostgreSQL Performance Tuning and PostgreSQL Security Training. Participants finish prepared to design, evaluate, protect, and deploy reliable intelligent applications using PostgreSQL with AI in operational environments with measurable business value.

Target Audience:

- PostgreSQL database administrators
- Database developers and database engineers
- AI and machine learning engineers
- Backend and full-stack developers
- Data engineers and data architects
- Software and solutions architects
- DevOps and cloud engineers
- Application integration specialists
- Business intelligence developers
- Information security professionals
- Technical team leaders
- Professionals developing RAG systems and AI chatbots



Targeted Organizational Departments:

- Database administration and data management
- Artificial intelligence and machine learning
- Software and application development
- Data engineering and analytics
- Enterprise and solutions architecture
- Cloud infrastructure and DevOps
- Cybersecurity and information governance
- Digital transformation and innovation
- Business intelligence and reporting
- Knowledge management

Targeted Industries:

- Banking, financial services, and insurance
- Government and public-sector organizations
- Telecommunications
- Healthcare and life sciences
- Technology and software development
- Energy, utilities, and infrastructure
- Retail and e-commerce
- Education and scientific research
- Legal and professional services
- Manufacturing and logistics
- Media and digital publishing

Course Offerings:

By the end of this course, participants will be able to:

- Configure and administer PostgreSQL for AI-enabled applications
- Apply AI-assisted SQL techniques while validating generated queries
- Install and configure the pgvector extension
- Generate and store PostgreSQL embeddings with relational data
- Use cosine, Euclidean, inner-product, and other distance measures
- Build exact and approximate vector similarity search queries
- Configure HNSW and IVFFlat indexes
- Combine vector search with metadata filters and relational queries
- Implement PostgreSQL Semantic Search

- Connect PostgreSQL to OpenAI and other language models

Training Methodology:

Training combines instructor-led explanations, demonstrations, laboratories, case studies, group design work, troubleshooting exercises, and feedback. Participants first practise PostgreSQL administration and AI-assisted SQL Training, checking generated queries for correctness, efficiency, and unsafe behaviour. Progressive pgvector Training laboratories cover embedding storage, distance operators, exact search, HNSW, IVFFlat, metadata filters, and query-plan analysis. Teams then map document ingestion, chunking, PostgreSQL embeddings, retrieval, augmentation, generation, and citation flows for a realistic PostgreSQL RAG Training scenario. Security exercises examine prompt injection, sensitive-information disclosure, poisoned vector content, excessive permissions, and insecure model outputs. Performance experiments compare recall, latency, storage, and index settings. Throughout the course, participants receive observable results and instructor feedback. The final activity integrates database design, semantic retrieval, chatbot logic, risk controls, evaluation criteria, monitoring, and deployment into one production-oriented intelligent application design.

Course Toolbox:

- PostgreSQL AI application architecture examples
- Database and vector-schema examples
- AI-assisted SQL prompt examples
- pgvector installation and configuration guidance
- Embedding-generation workflow examples
- Vector-distance operator reference
- HNSW and IVFFlat configuration examples
- Document ingestion and chunking checklist
- Semantic search design worksheet
- RAG pipeline architecture template

The course provides insights, demonstrations, templates, and examples of relevant tools. PostgreSQL environments, AI subscriptions, API credits, cloud services, software licences, hardware, and third-party tools are **NOT** provided.

Course Agenda:

Day 1: PostgreSQL Foundations and AI-Assisted SQL

- **Topic 1:** PostgreSQL architecture, databases, schemas, tables, and data types
- **Topic 2:** PostgreSQL installation, configuration, and server administration
- **Topic 3:** Tables, relationships, constraints, transactions, and conventional indexes
- **Topic 4:** SQL development for structured application data
- **Topic 5:** AI-assisted SQL generation, explanation, and optimization
- **Topic 6:** Validating AI-generated queries for accuracy, security, and performance
- **Reflection & Review:** Reviewing PostgreSQL fundamentals and correcting an AI-generated query

Day 2: PostgreSQL Embeddings and Vector Search

- **Topic 1:** Vector embeddings and semantic representation
- **Topic 2:** PostgreSQL as a relational and vector database
- **Topic 3:** Installing and configuring the pgvector extension
- **Topic 4:** Creating vector columns and storing PostgreSQL embeddings
- **Topic 5:** Cosine, Euclidean, inner-product, L1, Hamming, and Jaccard distances
- **Topic 6:** Building vector similarity search queries with metadata filters
- **Reflection & Review:** Comparing keyword, full-text, and PostgreSQL Semantic Search results

Day 3: Vector Indexing and RAG Development

- **Topic 1:** Exact and approximate nearest-neighbour search
- **Topic 2:** HNSW architecture, construction, search, and tuning parameters
- **Topic 3:** IVFFlat lists, probes, recall, and query-performance trade-offs
- **Topic 4:** Document extraction, cleaning, segmentation, and chunking
- **Topic 5:** Embedding pipelines and PostgreSQL vector indexing
- **Topic 6:** Retrieval, augmentation, generation, and source grounding
- **Reflection & Review:** Designing a PostgreSQL RAG pipeline for enterprise documents



Day 4: PostgreSQL Chatbots and Secure AI Integration

- **Topic 1:** Connecting PostgreSQL to OpenAI and other language models
- **Topic 2:** Prompt construction using retrieved database context
- **Topic 3:** Developing conversational retrieval and database-aware chatbots
- **Topic 4:** Managing conversation history, citations, and source attribution
- **Topic 5:** Controlling roles, privileges, row access, credentials, and sensitive data
- **Topic 6:** Preventing prompt injection, poisoned embeddings, unauthorized retrieval, and data leakage
- **Reflection & Review:** Reviewing and strengthening a PostgreSQL chatbot security architecture

Day 5: Performance, Evaluation and Deployment

- **Topic 1:** Measuring retrieval relevance, recall, precision, latency, and groundedness
- **Topic 2:** Using EXPLAIN and execution plans for vector-query analysis
- **Topic 3:** Tuning memory, connections, storage, queries, and vector indexes
- **Topic 4:** Monitoring, logging, backup, recovery, and lifecycle management
- **Topic 5:** Deploying PostgreSQL AI applications on-premises and in the cloud
- **Topic 6:** Designing an end-to-end intelligent application with PostgreSQL
- **Reflection & Review:** Presenting and reviewing a secure, production-ready RAG application design

FAQ:

What specific qualifications or prerequisites are needed for participants before enrolling in the course?

Participants should understand basic database concepts and have introductory SQL knowledge. Previous experience with PostgreSQL, Python, APIs, cloud platforms, or AI applications is helpful but not mandatory. The technical exercises can be adjusted to participants' development and database experience.

How long is each day's session, and is there a total number of hours required for the entire course?

Each day's session is generally structured to last around 4-5 hours, with breaks and interactive activities included. The total course duration spans five days, approximately 20-25 hours of instruction.

Does PostgreSQL replace the need for a dedicated vector database?

PostgreSQL can store, index, filter, and search embeddings through pgvector while retaining SQL, transactions, joins, metadata, access controls, backup, and replication capabilities. This makes it suitable for many semantic search and RAG applications. A separate vector database may still be considered when an application has exceptional distribution, scale, latency, or specialized operational requirements.

How This Course is Different from Other PostgreSQL AI Training Courses:

Unlike a general Vector Database Course or isolated chatbot workshop, this programme treats PostgreSQL as the operational foundation of the complete AI application lifecycle. It connects relational modelling, administration, AI-assisted SQL Training, PostgreSQL embeddings, pgvector indexing, semantic retrieval, RAG orchestration, security, performance, and deployment in one coherent learning journey. Participants compare exact search with HNSW and IVFFlat, examine recall-and-latency trade-offs, and combine vector similarity search with metadata filters and PostgreSQL access controls. The PostgreSQL RAG Training component goes beyond a simple demonstration by addressing ingestion quality, chunking, grounding, citations, retrieval evaluation, prompt injection, sensitive-information exposure, and vector or embedding weaknesses. It also draws on PostgreSQL documentation, foundational RAG research, OWASP guidance, and the NIST Generative AI risk profile. Consequently, participants learn not only how to build an intelligent application, but how to test, govern, tune, monitor, and deploy it responsibly in an enterprise environment under realistic operational and governance constraints.