multisoft
Systems

SAP HANA Vector Engine & GenAI

COURSE CONTENT

GET IN TOUCH

Multisoft Systems
B - 125, Sector - 2, Noida

(+91) 9810-306-956

info@multosoftsystems.com

www.multosoftsystems.com

About Multisoft

Train yourself with the best and develop valuable in-demand skills with Multisoft Systems. A leading certification training provider, Multisoft collaborates with leading technologies to bring world-class one-on-one and certification training. Intending to empower professionals and businesses across the globe, we offer more than 1500 training courses, which are delivered by Multisoft's global subject matter experts. We offer tailored corporate training, a comprehensive learning solution with lifetime e-learning access, after-training support and globally recognized training certificates.

About Course

The SAP HANA Vector Engine & GenAI course focuses on combining enterprise data with Generative AI using SAP HANA Cloud. It covers vector embeddings, similarity search, semantic retrieval, vector data management, RAG workflows, LLM integration, and context-aware AI applications. Participants gain practical understanding of using SAP HANA Vector Engine as a retrieval layer for intelligent search, enterprise knowledge assistants, and other GenAI-driven business scenarios.

Module 1: SAP HANA Cloud and Generative AI Fundamentals

- ✓ SAP HANA Cloud architecture overview
- ✓ Role of SAP HANA in modern data architectures
- ✓ Introduction to Generative AI
- ✓ Large Language Models and enterprise applications
- ✓ Structured versus unstructured enterprise data
- ✓ Limitations of traditional keyword-based retrieval
- ✓ Role of semantic search in GenAI
- ✓ SAP ecosystem for enterprise AI
- ✓ Position of SAP HANA Vector Engine in GenAI architecture

Module 2: SAP HANA Vector Engine Architecture

- ✓ Introduction to SAP HANA Cloud Vector Engine
- ✓ Vector processing concepts
- ✓ Vector Engine components and capabilities
- ✓ Vector processing within SAP HANA Cloud
- ✓ Combining relational and vector data
- ✓ Vector workloads and enterprise data
- ✓ Vector Engine architecture patterns
- ✓ Data flow across GenAI applications
- ✓ Understanding vector-based retrieval architecture

Module 3: Vector Embeddings and Vector Data Fundamentals

- ✓ Understanding vectors and multidimensional representations
- ✓ What are embeddings?
- ✓ Text-to-vector transformation
- ✓ Semantic representation of enterprise content
- ✓ Embedding dimensions
- ✓ Embedding spaces
- ✓ Similarity between embeddings
- ✓ Dense vector concepts
- ✓ Documents, chunks, and embeddings
- ✓ Embedding model selection considerations
- ✓ Embedding lifecycle in GenAI applications

Module 4: Vector Data Management in SAP HANA

- ✓ Vector data representation in SAP HANA
- ✓ REAL_VECTOR data type
- ✓ Creating tables for vector data
- ✓ Storing embeddings with business data
- ✓ Loading vector information
- ✓ Converting vector representations
- ✓ Vector-related SQL operations
- ✓ Managing vector columns
- ✓ Combining vector and conventional attributes
- ✓ Querying structured and vector information

- ✓ Managing document metadata with embeddings

Module 5: Vector Similarity Search and Retrieval

- ✓ Fundamentals of similarity search
- ✓ Semantic similarity concepts
- ✓ Query vectors
- ✓ Distance and similarity measurements
- ✓ Cosine similarity concepts
- ✓ Euclidean distance concepts
- ✓ Dot-product-based similarity
- ✓ Vector distance functions
- ✓ Ranking retrieved results
- ✓ Top-K retrieval
- ✓ Threshold-based retrieval
- ✓ Metadata-assisted filtering
- ✓ Hybrid retrieval patterns
- ✓ Evaluating retrieval relevance

Module 6: Embedding Models and LLM Integration

- ✓ Role of embedding models
- ✓ Embedding generation workflow
- ✓ Query and document embeddings
- ✓ Choosing appropriate embedding models
- ✓ Understanding LLM context
- ✓ Prompt and context relationships

- ✓ Connecting enterprise information with LLMs
- ✓ Model endpoint considerations
- ✓ Embedding consistency
- ✓ Handling model changes and re-embedding
- ✓ Token and context-window considerations
- ✓ Enterprise LLM integration patterns

Module 7: Semantic Search with SAP HANA Vector Engine

- ✓ Traditional search versus semantic search
- ✓ Semantic retrieval workflow
- ✓ Preparing enterprise documents
- ✓ Document extraction and preprocessing
- ✓ Chunking strategies
- ✓ Generating embeddings for chunks
- ✓ Storing chunks and metadata
- ✓ Creating semantic queries
- ✓ Executing vector similarity searches
- ✓ Ranking relevant passages
- ✓ Combining semantic search with business filters
- ✓ Improving retrieval quality
- ✓ Semantic enterprise knowledge search scenario

Module 8: Retrieval-Augmented Generation (RAG)

- ✓ Introduction to RAG
- ✓ Why RAG is used with LLMs

- ✓ RAG architecture components
- ✓ Ingestion pipeline
- ✓ Document preprocessing
- ✓ Chunking and embedding pipeline
- ✓ Vector storage layer
- ✓ User-query embedding
- ✓ Retrieval workflow
- ✓ Context augmentation
- ✓ Prompt construction
- ✓ LLM response generation
- ✓ Grounding responses with enterprise data
- ✓ Reducing hallucination risks
- ✓ RAG versus model fine-tuning
- ✓ RAG architecture with SAP HANA Vector Engine

Module 9: Building RAG Applications with SAP HANA

- ✓ Designing an SAP HANA-based RAG pipeline
- ✓ Preparing enterprise knowledge sources
- ✓ Data ingestion strategies
- ✓ Document chunk management
- ✓ Embedding generation
- ✓ Persisting embeddings in SAP HANA
- ✓ Performing vector retrieval
- ✓ Applying metadata filters
- ✓ Selecting relevant context
- ✓ Prompt enrichment

- ✓ Passing retrieved context to an LLM
- ✓ Processing generated responses
- ✓ Source and context traceability
- ✓ Retrieval quality evaluation
- ✓ End-to-end RAG application workflow

Module 10: SAP AI Core and Generative AI Hub Integration

- ✓ SAP AI architecture overview
- ✓ Introduction to SAP AI Core
- ✓ Role of SAP AI Launchpad
- ✓ Introduction to Generative AI Hub
- ✓ Model access through Generative AI Hub
- ✓ Foundation model concepts
- ✓ Model selection considerations
- ✓ Connecting GenAI models with SAP applications
- ✓ SAP HANA Vector Engine as a retrieval layer
- ✓ Embedding and inference workflows
- ✓ Prompt orchestration concepts
- ✓ Combining HANA data with Generative AI Hub models
- ✓ Enterprise GenAI integration architecture

Module 11: GenAI Application Development and Integration

- ✓ Designing enterprise GenAI applications
- ✓ Application architecture considerations
- ✓ Backend integration with SAP HANA Cloud

- ✓ Connecting applications to vector search
- ✓ Query processing workflow
- ✓ Context retrieval and augmentation
- ✓ LLM request and response handling
- ✓ API-based integration concepts
- ✓ Integration with SAP BTP applications
- ✓ CAP-based application integration concepts
- ✓ Conversational application patterns
- ✓ Knowledge assistants and enterprise copilots
- ✓ Handling sessions and conversational context
- ✓ Error handling and fallback strategies

Module 12: Security, Governance, Performance, and Production Readiness

- ✓ Enterprise AI security considerations
- ✓ SAP HANA authentication and authorization concepts
- ✓ Protecting sensitive enterprise information
- ✓ Role-based data access
- ✓ Securing model and application connections
- ✓ Data privacy considerations
- ✓ Responsible AI principles
- ✓ Prompt and response governance
- ✓ Grounding and information traceability
- ✓ Vector storage considerations
- ✓ Retrieval performance optimization

- ✓ Managing large vector datasets
- ✓ Query optimization considerations
- ✓ Monitoring GenAI applications
- ✓ Cost and resource considerations
- ✓ Production deployment considerations

Module 13: Enterprise GenAI Use Cases and Capstone Project

Enterprise Use Cases

- ✓ Intelligent enterprise knowledge search
- ✓ Product information assistant
- ✓ Customer support knowledge assistant
- ✓ HR policy and employee information assistant
- ✓ Technical documentation search
- ✓ Procurement and supplier information retrieval
- ✓ Financial document intelligence
- ✓ SAP business data question-answering scenarios
- ✓ Context-aware enterprise copilots

Capstone Project – SAP HANA Vector Engine RAG Application

- ✓ Define enterprise GenAI use case
- ✓ Prepare sample knowledge documents
- ✓ Design document chunking strategy
- ✓ Generate vector embeddings
- ✓ Store vectors and metadata in SAP HANA

- ✓ Implement similarity-based retrieval
- ✓ Apply business and metadata filters
- ✓ Retrieve relevant contextual information
- ✓ Construct augmented prompts
- ✓ Integrate an LLM
- ✓ Generate context-grounded responses
- ✓ Test retrieval accuracy
- ✓ Evaluate generated responses
- ✓ Review security and governance considerations
- ✓ Present complete GenAI solution architecture