JUPITER PRIME
Consulting • Training • Advisory
جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Enterprise Search and Knowledge Retrieval with AI Training Course

13 – 17 December 2026
Manama

Enterprise Search and Knowledge Retrieval with AI Training Course

Ref: 774_29255 **Date:** 13 – 17 December 2026 **Location:** Manama **Fees:** 4900 GBP

Enterprise Search and Knowledge Retrieval with AI is a 1-week training course for knowledge managers, information architects and IT leads who build a retrieval blueprint for an answer assistant that serves staff from company policies, manuals and records. Generative models now sit on top of document stores, and an assistant that cites a wrong version damages trust. Jupiter Prime Training runs this course on enterprise search and knowledge retrieval with AI.

Why Enterprise Search and Knowledge Retrieval with AI Matters Now

Three forces meet in this subject. Large language models made conversational answers expected, content estates sprawl across intranets and drives, and security teams ask who may see which sensitive passage. An organisation that grounds its assistant in curated, permission-aware sources gives staff fast answers with citations; one that indexes everything blindly spreads outdated procedures and exposes confidential material. The course rests on ISO 30401 Knowledge management systems and proven practice in information retrieval, semantic search and retrieval augmented generation RAG evaluation.

Who Gains Most from Enterprise Search and Knowledge Retrieval with AI

- Those who decide which collections an assistant may answer from
- Those who design metadata schemes, content types and navigation across document libraries
- Those who run search platforms, connectors and indexing pipelines
- Those who set access rules for sensitive records and audit how they are applied
- Those who measure whether answers resolve staff questions

A cross-functional group of knowledge, architecture and platform leads from one organisation gains most, since assistants fail where content, permissions and infrastructure meet.

Enterprise Search and Knowledge Retrieval with AI Capabilities

Participants judge where keywords suffice, where vectors help and where a knowledge graph earns its upkeep.

- Assess a content estate for freshness, duplication and ownership — the decision on which sources to index
- Design a chunking and metadata scheme for policies and manuals — the information architect and platform team
- Specify a hybrid retrieval pipeline with keyword, embedding and reranking stages — the IT lead choosing components
- Apply document-level and passage-level access control to answers — the security owner and auditors
- Build a test set of real staff questions with expected sources — the release decision for the assistant
- Communicate answer quality, citation accuracy and gaps to sponsors — the knowledge strategy owner

Key Enterprise Search and Knowledge Retrieval with AI Terms

- Retrieval augmented generation RAG — a language model writing its answer from passages fetched from a trusted collection.
- Embedding — numbers representing a text's meaning so similar passages sit close together.
- Vector index — a store returning passages whose embeddings lie nearest a question.
- Hybrid search — retrieval combining keyword matching with vector similarity.
- Knowledge graph — a network of entities such as products, sites and policies and the links between them.
- Grounding — the link between each statement in an answer and the source passage that supports it.

Enterprise Search and Knowledge Retrieval with AI Day by Day

Day 1: Content Estates and Search Behaviour

- Search Log Analysis — Which Questions Staff Ask Most Often
- Content Inventory and Ownership Map — Which Collections Are Fit to Index
- Version and Retention Rules — Removing Superseded Policies Before Indexing
- Lexical Search Baseline with BM25 — Where Keywords Already Suffice
- ISO 30401 Knowledge Context — Linking Retrieval to Knowledge Strategy

Day 2: Embeddings, Chunking and Hybrid Retrieval

- Chunking by Heading and Clause — Passage Size for Precise Citation
- Embedding Model Selection — Language Coverage, Cost and Domain Fit
- Vector Index Configuration — Recall Versus Latency Trade-Offs
- Hybrid Ranking and Rerankers — Ordering Candidate Passages for the Model
- Metadata Filters by Department and Document Type — Narrowing Scope

Day 3: Knowledge Graphs and Answer Generation

- Entity Extraction from Manuals and Contracts — Building Graph Nodes
- Ontology and Controlled Vocabulary — Consistent Names Across Collections
- Graph-Assisted Retrieval — Answering Multi-Hop Questions About Assets
- Prompt Templates with Citations — Grounded Answers Staff Can Verify
- Refusal and Escalation Rules — When the Assistant Hands Over to Experts

Day 4: Permissions, Evaluation and Operations

- Permission Trimming at Query Time — Keeping Restricted Passages Hidden
- Golden Question Sets — Measuring Retrieval Recall and Answer Faithfulness
- Hallucination and Citation Audits — Deciding Release Readiness
- Feedback Capture and Content Gap Reports — Prioritising New Articles
- Index Refresh and Connector Monitoring — Keeping Answers Current

Day 5: Practising the Retrieval Blueprint

- Use Case Selection for an HR Policy or Maintenance Assistant — Scoping the Pilot
- Source Curation and Chunking Plan — Applying Day Two Choices
- Access Model and Test Set — Applying Day Four Controls
- Sponsor Briefing on Risks and Measures — Securing Agreement to Proceed
- Retrieval Blueprint for an Answer Assistant — Completing the Team Output

Applying Enterprise Search and Knowledge Retrieval with AI After the Course

- Review search logs and the inventory of one busy collection, retiring superseded documents
- Build a golden question set from real staff queries and score the current search tool as a baseline

Enterprise Search and Knowledge Retrieval with AI Questions FAQ

What is enterprise search and knowledge retrieval with AI, and why does it matter now?

It is the practice of finding and assembling trusted passages from company content so an AI assistant can answer with citations. It matters because ungrounded assistants spread outdated or restricted information.

Which roles gain most from enterprise search and knowledge retrieval with AI?

Knowledge managers, information architects and IT leads gain most. You decide which sources are indexed, secured and tested before an assistant reaches staff in sectors such as energy or healthcare.

How does enterprise search and knowledge retrieval with AI differ from a general AI course?

It focuses on retrieval over company documents rather than on AI strategy or model building. You work on chunking, embeddings, ranking and permission trimming.

What do you bring back from enterprise search and knowledge retrieval with AI?

You bring back a retrieval blueprint for one answer assistant. It sets out curated sources, chunking, the access model and the measures for approving release.

Enterprise Search and Knowledge Retrieval with AI: The Way Forward

The retrieval blueprint for an answer assistant first informs which collection goes live and under which access rules. Knowledge managers, information architects and IT leads use it to curate sources, label passages and configure the retrieval pipeline. After release, revisit the golden question set, since staff questions shift faster than documents.