JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Generative AI Knowledge Assistants and RAG Pilots Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Generative AI Knowledge Assistants and RAG Pilots Training Course

Generative AI Knowledge Assistants and RAG Pilots is a 1-week training course for operations managers, analysts and project leads who need an assistant that answers from their own manuals and policies, ending with a RAG pilot charter. Organisations are replacing open chat tools with assistants that cite internal sources, and the gap between a demo and a trusted answer is where effort is lost. Jupiter Prime Training runs this course for teams testing that gap.

Why Generative AI Knowledge Assistants and RAG Pilots Matter Now

Three forces shape this field. Staff already paste internal text into public chat tools, so leaders need a sanctioned alternative. Language models answer fluently yet invent facts, which makes retrieval-augmented generation RAG, with answers tied to cited passages, the practical route for policy and procedure questions. Auditors and risk owners also expect evidence of control. An organisation that grounds its assistant well shortens search time and reduces rework; one that does not spreads confident errors. The course rests on ISO/IEC 42001, the AI management system standard, and the NIST AI Risk Management Framework NIST AI RMF.

Who Gains Most from Generative AI Knowledge Assistants and RAG Pilots

- Operations owners responsible for procedures and frontline queries
- Analysts who prepare and tag searchable document collections
- Project leads who scope, schedule and report on an assistant trial
- Knowledge and records custodians who approve current sources
- Risk and quality coordinators who sign off answer accuracy

A cross-functional trial squad gains most when its operations owner, analyst and risk coordinator attend together.

Generative AI Knowledge Assistants and RAG Pilots Capabilities

Participants judge where a grounded assistant earns its place and how to prove its answers before launch.

- Assess which recurring questions suit a retrieval-backed assistant — use-case owners choosing a first scope
- Specify a source collection with chunking and metadata rules — records custodians approving content
- Design prompt templates that force citations and refusals — operations owners setting answer standards
- Build a golden question set and score faithfulness — quality coordinators accepting results
- Apply access, logging and escalation controls drawn from ISO/IEC 42001 — risk owners
- Communicate trial findings with a go, adjust or stop recommendation — sponsors funding scale-up

Key Generative AI Knowledge Assistants and RAG Pilots Terms

- Retrieval-augmented generation — a pattern in which a model answers using passages fetched from an approved document store rather than memory alone.
- Chunking — splitting long documents into short passages for precise search.
- Embedding — a numeric representation of text meaning that lets a vector index find passages similar to a question.
- Grounding — tying each statement in an answer to a cited source passage.
- Hallucination — a fluent model statement that no supplied source supports.
- Golden question set — a fixed list of real questions with agreed correct answers.

Generative AI Knowledge Assistants and RAG Pilots Day by Day

Day 1: Where Grounded Assistants Fit

- Knowledge Question Inventory — Which Recurring Queries Justify an Assistant
- Open Chat Versus Retrieval Patterns — Choosing the Right Answer Architecture
- Value and Effort Grid — Ranking Candidate Assistant Scopes
- Stakeholder Map for Document Owners — Securing Source Approval Early
- NIST AI RMF Map Function — Framing Context and Intended Use

Day 2: Building the Source Collection

- Document Currency Audit — Removing Superseded Policies Before Indexing
- Chunking and Metadata Rules — Improving Passage Relevance and Filtering
- Embeddings and Vector Index Basics — Setting Search Expectations Realistically
- Access Tiers and Permission Mirroring — Keeping Restricted Files Restricted
- Retrieval Precision Spot Checks — Deciding When Sources Need Rework

Day 3: Prompt Patterns and Answer Design

- System Prompt With Citation Rules — Forcing Traceable Answers
- Refusal and Fallback Wording — Handling Questions Outside the Collection
- Few-Shot Answer Examples — Shaping Tone for Frontline Users
- Structured Output Templates — Feeding Answers Into Tickets and Forms
- Human Escalation Triggers — Routing Sensitive Queries to Specialists

Day 4: Proving Answer Quality and Control

- Golden Question Set Construction — Agreeing What Correct Looks Like
- Faithfulness and Hallucination Scoring — Accepting or Rejecting Releases
- User Feedback Loops and Logging — Spotting Drift After Launch
- ISO/IEC 42001 Control Mapping — Evidencing Oversight to Auditors
- Prompt Injection and Data Leakage Tests — Protecting Confidential Content

Day 5: Practising the RAG Pilot Charter

- Scope Statement From the Question Inventory — Fixing Trial Boundaries
- Source and Access Plan Rehearsal — Confirming Owners and Permissions
- Evaluation Scorecard Dry Run — Setting Go and Stop Thresholds
- Trial Budget, Roles and Timeline Sketch — Securing Sponsor Agreement
- RAG Pilot Charter Completion — Presenting a Ready-to-Launch Trial

Applying Generative AI Knowledge Assistants and RAG Pilots After the Course

- Run a document currency audit on one procedure library, tagging owner, version and access tier before indexing.
- Draft a golden question set of real frontline queries and score the first assistant build against it.

Generative AI Knowledge Assistants and RAG Pilots Questions FAQ

What are Generative AI Knowledge Assistants and RAG Pilots, and why do they matter now?

They are assistants that answer from approved internal documents, tested first in a limited trial. They matter now because staff already use public chat tools, and cited answers reduce invented facts in policy and service queries.

Which roles gain most from Generative AI Knowledge Assistants and RAG Pilots?

Operations owners, analysts and project leads gain most. Records custodians and quality coordinators also benefit, since you decide which sources are searched and how answers are scored.

How does Generative AI Knowledge Assistants and RAG Pilots differ from a general applied AI course?

It focuses on one pattern: retrieval-backed answers from your own documents. A general applied AI course surveys many tools; this one goes deep on source preparation, citation prompts and faithfulness scoring.

What do you bring back from Generative AI Knowledge Assistants and RAG Pilots?

You bring back a trial charter for one assistant. It sets scope, approved sources, access rules, prompt patterns and the thresholds that decide whether the trial expands or stops.

Generative AI Knowledge Assistants and RAG Pilots: The Way Forward

The RAG pilot charter is the starting point. Its first use is the sponsor's decision on which document collection to index and which questions the assistant may answer. Operations owners and risk coordinators use it to run the trial. After the first scored release, revisit the golden questions, since real user queries usually reveal gaps the original list missed.