



Knowledge Graphs and Ontology Engineering

Course: RDF, OWL, SPARQL and GraphRAG



Knowledge Graphs and Ontology Engineering Course: RDF, OWL, SPARQL and GraphRAG

Introduction:

Knowledge graphs and ontology engineering with RDF, OWL, SPARQL and GraphRAG is a 5-day course for data architects, semantic modellers, data engineers and AI retrieval teams, ending with a Knowledge Graph Blueprint for a case organisation. Many organisations hold one citizen, patient or customer in disconnected tables, so cross-system questions need manual joins and language-model assistants answer without grounded context. Nominees already design schemas, pipelines or analytics; the course turns that experience into triples, classes, shapes and graph queries through hands-on labs on RDF stores and a property graph database. CoreConcept Training Center delivers this knowledge graphs course.

Course Objectives:

- Model a business domain as classes, properties and individuals in RDF and RDFS, scoped by competency questions
- Encode ontology axioms in OWL 2 and select the EL, QL or RL profile that suits the reasoning load
- Write SPARQL queries that traverse, federate and aggregate linked data held in several sources
- Validate graph data with SHACL shapes and organise controlled vocabularies as SKOS concept schemes
- Build and query a labelled property graph in Neo4j with Cypher and align duplicate entities into single nodes
- Design a GraphRAG retrieval flow and prepare a Knowledge Graph Blueprint covering ontology, pilot scope, ingestion and query plan

Target Audience:

- Data architecture teams responsible for enterprise information models and integration schemas
 - Data engineering staff who build ingestion pipelines and maintain graph or relational stores
 - Information and taxonomy staff who curate vocabularies, classification schemes and search metadata
 - AI and analytics engineering staff who connect language models and machine learning to enterprise data
 - Open data and digital service staff in public bodies who publish and link government datasets
-

Course Outline:

Day 1: Graph Thinking, Triples and Domain Scoping

- Knowledge Graph Anatomy of Nodes, Edges and Literals
- RDF Triple Model With IRIs, Blank Nodes and Datatypes
- Turtle and JSON-LD Serialisation of Sample Government Records
- Competency Questions for Scoping a Graph Use Case
- Current Data Landscape Mapping to Candidate Graph Entities

Day 2: Ontology Languages, Reasoning and Vocabulary Standards

- RDFS Class Hierarchies, Domains and Ranges in Practice
- OWL 2 Axioms for Equivalence, Disjointness and Cardinality
- OWL 2 EL, QL and RL Profile Selection
- TBox, ABox and RBox Separation in Ontology Design
- SKOS Concept Schemes With Broader, Narrower and Related Links

Day 3: Querying, Validation and Property Graph Labs

- SPARQL SELECT, CONSTRUCT and Property Path Query Lab
- SPARQL Federation Across Linked Open Data Endpoints
- SHACL Shapes Graph Authoring for Data Graph Validation
- Neo4j Labelled Property Graph Modelling and Cypher Queries
- RDF Versus Property Graph Trade-Offs Including ISO GQL

Day 4: Entity Alignment, Graph Analytics and GraphRAG Risks

- Entity Alignment Across Sources Using Shared Attributes and Substructures
- Graph Embeddings for Link Prediction and Similarity Search
- GraphRAG Pipeline Linking Language Model Retrieval to Graphs
- Ontology Drift, Versioning and Reasoner Performance Troubleshooting
- Access Control and Provenance Tracking on Sensitive Graph Data

Day 5: Lab Build of the Knowledge Graph Blueprint

- Government Open Data Lab Using the DCAT Catalogue Vocabulary
 - Health Records Graph Lab Linking Patients, Encounters and Clinical Codes
 - Customer 360 Graph Lab Joining Accounts, Interactions and Products
 - Pilot Scope, Ingestion Pipeline and SPARQL Query Plan Drafting
 - Knowledge Graph Blueprint Completion and Peer Review Panel
-

Skills You Will Gain:

- RDF Data Modelling
- OWL Ontology Design
- SPARQL Query Writing
- SHACL Data Validation
- Taxonomy and Thesaurus Design
- Property Graph Modelling
- Graph Entity Alignment
- GraphRAG Retrieval Design

Why Attend This Course:

- Deliver a Knowledge Graph Blueprint, covering ontology, pilot scope, ingestion and query plan, to the data architecture board and the sponsoring business unit
- Choose between an RDF store and a property graph database for a given use case, and between OWL 2 profiles for its reasoning needs
- Avoid ontology rework, unvalidated graph data and ungrounded language-model answers that erode trust in AI services
- Coach data engineers and analysts in writing SPARQL and Cypher queries and in applying the shared vocabulary

Conclusion:

Back at work, the participant hands the data architecture board and the sponsoring business unit a Knowledge Graph Blueprint for one priority domain, such as citizen services, patient pathways or customer 360. The board uses it to approve the pilot scope, the choice between an RDF store and a property graph, and the sources to ingest first. After the pilot's first release, the unit should review query response times, SHACL validation results and the accuracy of GraphRAG answers against the competency questions, then revise the ontology and the next ingestion wave.

Frequently Asked Questions FAQ:

What should participants know before a knowledge graphs and ontology engineering course?

Participants should be comfortable with relational data models, basic SQL and reading JSON or XML. No prior semantic web experience is needed, but bringing a sample dataset or data dictionary from their own domain makes the labs and the blueprint more relevant.

How does knowledge graphs and ontology engineering differ from a data management framework or master data course?

This course builds semantic graph models, ontologies and graph queries. Data management framework courses assess capability across many disciplines, and master data courses focus on quality rules and golden records in hubs; both touch graphs only briefly.

Why does ontology engineering matter for knowledge graphs used with GraphRAG?

An ontology gives a knowledge graph a shared schema, so a language model retrieves connected, typed facts rather than loose text fragments. GraphRAG uses those relationships to answer questions that span several documents or systems, with a traceable source for each statement.

What do participants take back to work from a knowledge graphs and ontology engineering course?

Participants take back a Knowledge Graph Blueprint for a case organisation, with an ontology sketch, SHACL shapes, pilot scope, ingestion sequence and SPARQL or Cypher query plan, plus lab query examples to reuse with their own teams.