# RAJESH KUMAR GUPTA

**ENTERPRISE ARCHITECT | AGENTIC AI | ONTOLOGY & KNOWLEDGE GRAPH | SEMANTIC DATA ARCHITECTURE**

Basking Ridge, NJ • Open to relocation anywhere in the U.S. • +1 732-421-2092 • rajesh21gupta@gmail.com

linkedin.com/in/rajeshgupt • github.com/aisemanticexpert

## EXECUTIVE PROFILE

Enterprise Architect and Semantic AI leader with 15+ years designing enterprise data, knowledge graph, ontology, and AI platforms across insurance, financial services, healthcare, and manufacturing. Deep expertise in enterprise and domain ontology, canonical data models, GraphRAG, agentic AI, semantic governance, and cloud-native data architecture. Hands-on with OWL/RDF/SHACL/SKOS, Neo4j/Neptune/MarkLogic, Snowflake, LangGraph/LangChain, AWS, Python, and Java, building governed AI architectures that connect structured and unstructured data to business-ready semantic services. Published ontology researcher (IEEE, Elsevier) and conference speaker with a track record of mentoring engineering teams and driving enterprise adoption of semantic standards.

## CORE ARCHITECTURE & TECHNOLOGY

**Architecture:** Enterprise Architecture, Semantic Architecture, Canonical Enterprise Data Models, Domain-Driven Design, API & Data Integration, Event & Streaming Architecture

**Semantic / KG:** Enterprise & Domain Ontology, OWL 2 DL, RDF/TTL, JSON-LD, SHACL, SKOS, SPARQL, Cypher, FIBO, ACORD, Protégé, RDFLib, Apache Jena, Neo4j, Amazon Neptune, MarkLogic, GraphDB

**AI / Retrieval:** Agentic AI, Multi-Agent Systems, GraphRAG, Hybrid Graph + Vector Retrieval, Model Context Protocol (MCP), LangGraph, LangChain, LlamaIndex, Claude/GPT, Text2Cypher/Text2SPARQL, context engineering, human-in-the-loop AI

**Data / Cloud:** Snowflake, Databricks/Spark, Kafka, Airflow, AWS (S3, Glue, Bedrock, SageMaker, Lambda, Neptune), Docker, Kubernetes, OpenShift, CI/CD

**Governance:** Ontology governance & versioning, data lineage, metadata management (DataHub, Collibra, Amundsen), data quality (Great Expectations), RBAC, AI governance, MiFID II, FINRA, HIPAA, FDA, EMA, GxP

**Domains:** Insurance (Policy, Claims, FNOL, Party, Product, Coverage, Risk, Underwriting, P&C, Life, Commercial, Consumer), Financial Services & Wealth Management, Healthcare & Life Sciences (HL7, FHIR, SNOMED CT, RxNorm, ICD-10/11, LOINC)

**Programming:** Python, Java, SQL, Cypher, SPARQL, TypeScript

## PROFESSIONAL EXPERIENCE

**Chubb — Enterprise Architect** **15 Jul 2026 – Present**

- Lead enterprise architecture for an ontology-centered semantic layer spanning Commercial, Consumer, and Life insurance, aligning core Policy, Claim, FNOL, Party, Product, Coverage, and Risk concepts across departmental canonical models.
- Define enterprise and domain ontology architecture using OWL 2, RDF/TTL, SHACL, and SKOS, with reusable modeling patterns, semantic constraints, controlled vocabularies, and alignment to insurance standards including ACORD and relevant FIBO concepts.
- Architect semantic data flows that map structured and unstructured source data through enterprise ontology and knowledge-graph layers into governed APIs, analytics, GraphRAG, LLM, and agentic-AI consumption patterns.
- Design FNOL and claims semantic patterns for entity/relation extraction, ontology population, validation, inference, and traceability, enabling consistent meaning from document ingestion through downstream AI reasoning.
- Establish architecture guardrails for ontology governance, versioning, stewardship, mapping, reasoning, and reuse; partner with business, data, platform, and AI teams to drive canonical-model conformance and enterprise adoption.

**Bessemer Trust — Data Architect — Ontology & Knowledge Graph** **Oct 2025 – Jul 2026**

*Somerset, NJ*

- Designed a firm-wide enterprise ontology across insurance, wealth planning, and investment domains, creating a shared semantic model across 15+ business capability areas and aligning with FIBO and ACORD.
- Built an end-to-end semantic document pipeline for PURE, Chubb, and Berkley One insurance documents, extracting policy structures, clauses, and underwriting attributes into a unified knowledge graph.
- Engineered Git-based ontology release management with semantic versioning, automated SHACL validation, CI/CD deployment to Amazon Neptune, and formal governance for change, stewardship, and backward compatibility.
- Developed LangGraph-based GraphRAG, Text2Cypher, and Text2SPARQL agents with schema-aware grounding and self-correction for natural-language access to enterprise knowledge.
- Architected an enterprise data lineage strategy across investment, operations, compliance, and client-reporting functions, mapping source-to-consumption paths through 50+ pipelines for regulatory traceability and impact analysis.
- Built an interactive ontology and knowledge-graph visualization tool with class hierarchies, property paths, and SPARQL/Cypher query interfaces, accelerating adoption among non-technical stakeholders and SMEs.

**Barclays — AVP, Data Architect & Knowledge Graph Engineering Lead** **May 2023 – Oct 2025**

*Whippany, NJ*

- Architected an enterprise Neo4j financial knowledge graph with 500+ entity types and 2M+ nodes, integrating market, reference, and transactional data as a semantic layer for investment intelligence.
- Pioneered GraphRAG processing 10TB+ daily of regulatory documents, research, and market commentary using LangChain/GPT-4, achieving 92% entity-extraction accuracy.
- Designed a multi-dimensional asset ontology, temporal versioning, decision-lineage graph, and thin-pointer architecture connecting petabyte-scale Snowflake data to graph analytics and regulated audit trails.
- Built the Snowflake data platform across 50+ market data sources (Bloomberg, Reuters, internal feeds) processing 5TB+ daily, orchestrated by 200+ Apache Airflow DAGs at 99.99% uptime.

- Established DataHub metadata/catalog patterns across 1M+ data assets and a Great Expectations data-quality framework with 500+ validation rules and automated anomaly detection.
- Mentored 12+ engineers on Cypher, graph algorithms, ontology, and GraphRAG; documented business logic and derivations supporting MiFID II and FINRA compliance.

**S&P Global — Senior Data Architect & Knowledge Graph Engineer** **Dec 2022 – May 2023**
*New York, NY*

- Built a Neo4j-based manufacturing intelligence graph over 1M+ documents and supplier/component relationships, with GraphRAG and ontology-driven classification achieving 85% extraction accuracy.
- Integrated Snowflake, Airflow, Kafka, and AWS with automated schema-drift detection, Amundsen metadata discovery, lineage, and quality validation; trained 20+ engineers and data scientists.

**Johnson & Johnson — Data Architect & Semantic Knowledge Graph Lead** **Dec 2019 – Dec 2022**
*USA*

- Architected a pharmaceutical knowledge platform with Neo4j and MarkLogic, integrating a 500K+ entity drug-discovery ontology and serving 500+ researchers.
- Integrated SNOMED CT, RxNorm, ICD-10/11, LOINC, and FHIR using RDF/OWL/TTL and built GraphRAG pipelines over 2M+ clinical/research documents using BioBERT/SciBERT/PubMedBERT.
- Implemented Snowflake data harmonization across 100+ clinical studies annually with temporal versioning, ontology-driven RBAC, lineage, and governed audit trails supporting FDA, EMA, and GxP requirements.
- Delivered a training curriculum on Cypher, SPARQL, graph traversal, GraphRAG, and semantic reasoning to 30+ pharmaceutical data scientists and engineers.

**Scania — Data Architect & Graph Analytics Engineer** **Jun 2018 – Nov 2019**
*Sweden*

- Built an enterprise Neo4j knowledge graph for 50,000+ connected vehicles with an OWL predictive-maintenance ontology reaching 90% parts-recommendation accuracy, plus Kafka streaming, Kubernetes deployment, and real-time graph analytics.

**AT&T — Data Architect & Knowledge Graph Engineer** **Mar 2014 – May 2018**
*India*

- Architected ontology-driven Neo4j/MarkLogic knowledge platforms for enterprise content analytics serving 50,000+ employees, with RDF/OWL/SPARQL, Airflow ETL on AWS processing 100,000+ daily transactions, metadata governance, business glossary, and fine-grained access control.

**Healthcare & Government Projects — Semantic Data Architect / Ontology Engineer** **Aug 2009 – Mar 2014**
*India*

- Developed healthcare ontologies and semantic search using SNOMED CT, ICD, RxNorm, and LOINC for clinical decision support; created ontology-merging research published through IEEE and the SafeOn road-safety ontology.

## SELECTED ARCHITECTURE & AI PORTFOLIO

- **SwarmCare:** Prototyped a multi-agent healthcare coordination platform using LangGraph, FHIR R4, and Neo4j, with ontology-grounded workflows for specialist, insurance, pharmacy, monitoring, and care-coordination use cases.
- **MCP & Agent Tools:** Built TypeScript and Python Model Context Protocol servers — including an 8-tool HELIOS/RAGHeat server — exposing graph retrieval, reasoning, and analytics capabilities as tools for LLM agents.
- **Hybrid KG + Vector RAG:** Designed POCs over 100,000+ insurance and wealth documents using Amazon Neptune (SPARQL), OpenSearch vector retrieval, embeddings, and LangGraph orchestration for schema-grounded enterprise question answering.

## RESEARCH, PUBLICATIONS & LEADERSHIP

- **RAGHeat:** Designed a GraphRAG framework combining graph attention and heat-diffusion reasoning with hybrid Neo4j + vector retrieval, reporting a 13.2% improvement in financial recommendation accuracy at sub-1.5s retrieval latency.
- **IEEE:** "An Instance-Based Methodology for Merging Domain Ontology" — ontology integration research with 50+ citations.
- **Elsevier / Transportation Research:** "Road Safety Ontology – SafeOn: Overview & Design" — OWL-based semantic framework.
- **In progress:** "RAGHeat: A Diffusion-Based GraphRAG Framework for Financial Recommendation."
- **Speaker:** NextGen Banking Summit 2024 — Knowledge Graphs and GenAI in Financial Services; ongoing mentor and trainer on ontology, Cypher, GraphRAG, and semantic architecture.
- **Current research interest:** Human-AI collaborative ontology engineering — constrained LLM ontology enrichment with reasoning, SHACL validation, provenance, expert review, and versioned backward compatibility.

## EDUCATION & CERTIFICATIONS

M.Tech, Software Engineering — National Institute of Technology (NIT), Allahabad, India, 2009 | B.E., Computer Science — Rajiv Gandhi Technological University, Bhopal, India, 2007

**Certifications:** Neo4j Certified Professional | Snowflake SnowPro Core | AWS Solutions Architect | AWS Data Analytics Specialty | LangChain AI Development