



Welcome to graphsummit



Agenda - Morning

09:30	Welcome Stephen Chin, Neo4j
09:35	Keynote: the knowledge layer for trustworthy AI Emil Eifrem, Neo4j
10:20	Connect, Don't Collect: How Inigo Is Governing Data for Enterprise AI with a Knowledge Layer Soundappan Soundarraj, Inigo Insurance
10:50	Morning break
11:10	Neo4j product vision & roadmap Ivan Zoratti & Michael Hunger, Neo4j
12:10	Building an AI-Native Data Platform: How AstraZeneca Uses Graph Technology to Deliver Context, Interoperability, and Scalable AI Russell Adams, Capgemini
12:30 - 13:30	Lunch

Agenda - Afternoon

- | | |
|-------|---|
| 13:10 | Demo: From alert to decision: Neo4j GraphAware Financial Crime Intelligence
Christophe Willemson, Neo4j |
| 13:30 | Building a Sovereign System of Intelligence at HSF Kramer
Raj Toora, HSF Kramer |
| 14:00 | Q&A panel
Soundappan Soundarraaj, Inigo Insurance
Raj Toora, HSF Kramer
Emil Eifrem & Ivan Zoratti, Neo4j |
| 14:20 | Afternoon break |

Choice of breakouts

Breakout track 1 - Executive track

14:50	The Enterprise Context Layer: Why AI Needs More Than Clean Data Matt Aslett, ISG
15:15	Unlocking a Century of Knowledge: How L'Oréal Uses Graph Technology to Make Its Unstructured Data Heritage AI-Ready Nicolas Guinard, SIA Partners Hussein Jouni, L'Oréal
15:25	The value of Neo4j Rupert Walter, KPMG James Radley, BT Group
15:55	Break
16:15	The AI inflection point and the knowledge layer Rob McManus & Jesus Barrasa, Neo4j
17:00	Networking reception

Breakout track 2 - Hands-on GenAI workshop: build an AI agent on your own knowledge graph

14:50	Hands-on workshop Kemi Elizabeth Ojogbede & the Neo4j team
15:55	Break
16:05	Hands-on workshop
17:00	Networking reception





graphsummit

The knowledge layer for
trustworthy AI

Emil Eifrem, Chief Executive Officer, Neo4j











The background is a dark navy blue. In the top right and bottom left corners, there are abstract, organic shapes in shades of teal and light blue, resembling liquid droplets or flowing paint.

Help the world
to make sense of data




The biggest change in data
in decades...

happened in the last 12 months



This changes **everything**





Customer Call Scenario

01. CALL INITIATION

The agent handles a customer call. They want to refinance. A competitor has quoted them 40 basis points better than our bank.

02. VALUE AT RISK

On the table is a \$412,000 mortgage, four other products, and a nineteen-year relationship.


03. DECISION WINDOW

The agent has about thirty seconds to decide whether to match, beat, or let the customer walk.

*What product can I offer this customer,
and at what price point?*

UNRESOLVED ENTITIES

Who is calling?



ORIGINATION



Charles Dolan
borrower_id 8842

SERVICING



C. Dolan
loan A-3391

COLLECTIONS



CALL CENTER



Charlie Dolan
contact 77120

COLLATERAL




RISK



What product can I offer this customer,
and at what price point?

STALE DATA

What do they earn?



RETAIL BANK



MORTGAGE



\$145,000
origination

CARDS



\$180,000
self-reported

WEALTH



\$310,000
KYC
verified last month

RETIREMENT



RISK



2019


2022

Last month

What product can I offer this customer,
and at what price point?

CONFLICTING DATA

What tier are they?



RETAIL BANK

MORTGAGE

CARDS

WEALTH

RETIREMENT

RISK



Standard
retail ops

Silver
card services

Gold
private bank

*Deposit
Balance*

*Annual
Spend*

Assets Managed

**AI in the enterprise is not constrained
by model intelligence...**

**... The bottleneck is giving the model
access to the right context**

Everyone's converging on the same problem



I need my agents to **act reliably at scale**
across my landscape of **disparate data sources**



Gartner formalizes “**AI Context Platforms**” as an emerging category, moving beyond RAG infrastructure toward a reusable enterprise layer for delivering governed context to AI agents.



Market Overview for AI Context Platforms

13 August 2026 - ID G00857698 - 13 min read

By: Kjell Carlsson, Afraz Jaffri, Christopher Long, Radu Miclaus

Initiatives: Technologies and Markets; Rethink Analytics and AI Through Agentic Transformation


AI context platforms construct, govern and deliver reusable context to AI agents at runtime. As enterprises scale production agents, leaders responsible for AI should implement these platforms to develop context into a first-class enterprise asset, rather than rebuild it separately for every agent.

Context?




2017

London



2017

London



2017 London

YouTube Premium interface showing a video player with a diagram illustrating a person's relationships and activities.

The diagram features a central profile of a person's head, with various icons and labels connected by arrows, representing different aspects of their life:

- MARRIED TO**: Connected to an icon of a couple.
- PLAYS**: Connected to an icon of a grand piano.
- WATCHES**: Connected to an icon of a television set.
- BORN IN**: Connected to an icon of a map of Taiwan.
- DRIVE**: Connected to an icon of a car.
- FATHER_OF**: Connected to an icon of a family (two adults and a child).
- WORKS_AT**: Connected to an icon of a network or globe.

The video player interface includes a search bar, a "Create" button, and a progress bar at the bottom. The video title is "GraphConnect Europe 2017 Opening Keynote – Emil Eifrem".

The world has now caught up



Genie Ontology: a living context graph for your business

PRODUCT June 16, 2026

Genie One and Genie Agents are powered by **Genie Ontology**, an **automatic context layer**. Genie Ontology automatically extracts snippets of knowledge from tables, queries, dashboards, pipelines, and connected apps, and **organizes that knowledge into a living graph** of how a company works and what the data inside *actually* means. Genie has context about where to look, what to trust, and how to answer in a way that reflects how the company actually uses its data. That includes metric **definitions, business terms, unique calculations, and the relationships between concepts**, metrics, tables, and teams.



May 25, 2026 / 13 min read CORE PLATFORM

Copy post link

Ontology-grounded Reasoning with Cortex Agents



Ontology

The central system for orchestrating decisions across Human+AI teams

The Palantir **Ontology** is an operational layer for the organization. The Ontology sits on top of the digital assets integrated into the Palantir platform (**datasets, virtual tables, and models**) and connects them to their real-world counterparts, ranging from physical assets like plants, equipment, and products to concepts like customer orders or financial transactions. In many settings, **the Ontology serves as a digital twin** of the organization, containing both the semantic elements (objects, properties, links) and kinetic elements (actions, functions, dynamic security) needed to enable use cases of all types.

Microsoft Fabric

The *ontology (preview)* item (part of the **Fabric IQ (preview) workload**) digitally represents the enterprise vocabulary and **semantic layer that unifies meaning across domains** and OneLake sources. It defines enterprise concepts as **entity types** (like *Customer*), **properties** (like a Customer's *name* and *email*), and **relationships** (like *Customer places Order*), while clarifying the constraints of these terms. After defining your ontology, bind the entity type definitions to real data, so downstream tools can share the same language. Both humans and AI agents can use this language for cross-domain reasoning and decision-ready actions.



Knowledge Graph: Powering intelligent and context-aware search

Send feedback

Knowledge Graph enhances Gemini Enterprise by linking data across three core pillars: people, content, and interactions. Knowledge Graph improves search quality and experience by **understanding the relationship** between different instances and entities in your corpus, providing deeper, context-aware search.



Amazon Web Services now introduces Context Ontology Accelerator

Posted on: Jul 31, 2026

Today, **AWS** announces Context Ontology Accelerator, an open source accelerator that helps organizations build an **ontology of their business — a machine-readable model** of their products, customers, policies, and the rules that define how they operate so AI agents make accurate, consistent, explainable, and auditable decisions. Context Ontology Accelerator connects to your structured and/or unstructured data sources, and uses AI to draft an ontology; your domain experts review, edit, and approve every element of it. The approved **ontology** is stored in a **knowledge graph** you own, expressed in open W3C standards so it works with any standards-based tooling, and any agent can consume it through the included Model Context Protocol (MCP) server.



THE REALLY GREAT NEWS

We know how to do this now



20 years of graphs

10 years of graphs and AI

2+ years of LLMs and graphs
in production



I need my agents to **act reliably at scale**
across my landscape of **disparate data sources**

1

Meaning

2

Reach

3

Learning



I need my agents to **act reliably at scale**
across my landscape of **disparate data sources**

1

Meaning


2

Reach

3

Learning


Meaning



```
if (status == 2) { // Suspended pending review
    order.holdFulfillment( );
    notifyReviewTeam(order);
}
```

Order ID	Status	Amount
100423	1	1,240
100424	2	3,580
100425	1	890
100426	3	2,120

Meaning





2= suspended pending review



Order ID	Status	Amount
100423	1	1,240
100424	2	3,580
100425	1	890
100426	3	2,120

Meaning



2= ...uh...
I got nuthin'



Order ID	Status	Amount
100423	1	1,240
100424	2	3,580
100425	1	890
100426	3	2,120

Ontologies to the rescue!



Post



lux
@lux



you have to be over 40 to even discuss ontologies

2:29 AM · Jul 26, 2026 · **532** Views



2



1



12




Relevant ▾

Meaning




Metaphysics IV
"what's the structure of
reality itself?"

300 BC



"an explicit specification of
a conceptualization"


1990s




"a digital twin that connects
an organization's raw data
to real-world business
concepts"

2000s

From GraphConnect NYC 2018



Ontologies



“The nouns and verbs of your business”

Really: the type system for your data

A Real-World Enterprise Ontology Example

Policy Ontology
Access & usage rules

Organization Ontology
Units, rules, ownership

Business Process Ontology
Tasks, activities, decision points

Domain Ontology
Business entities & relationships

Data Product Ontology
Links business concepts to systems

Technical Ontology
Systems, sources, data sets

A Real-World Enterprise Ontology Example

*Customers, Bank Accounts,
Transactions, Credit Cards, etc*

Domain Ontology
Business entities & relationships

*Oracle instance ABC123
Customer table, f_name column*

Technical Ontology
Systems, sources, data sets



I need my agents to **act reliably at scale**
across my landscape of **disparate data sources**

1

Meaning


2

Reach

3

Learning

Agent Maturity Curve Is Still in Early Days.




Source: Gartner IT Spending Estimates, Forrester, U.S. Bureau of Labor Statistics, Redpoint analysis.

Note: U.S. software and services estimates derived from Gartner spending data adjusted for Forrester's estimate of 41% U.S. market share. Operational and knowledge work payroll estimates derived from U.S. Bureau of Labor Statistics Occupational Employment and Wage Statistics (2024) across selected occupational groups based on Redpoint classification and judgment. TAM analysis includes US only.



©2026 Redpoint Ventures. Proprietary and confidential. Do not copy without permission.

I need a mortgage lender agent that drafts the best possible email or text message for a refinancing upsell offer.



Reach

What's the best refinancing message?



Agents in the enterprise require multi-hop reasoning

 Google Research

"Current single-step RAG systems **weren't designed for the multi-source, multi-hop queries of modern business workflows.**"

- Google Research, [June 2026](#)

ANTHROPIC

"You have a pile of unstructured documents and need to answer questions that span them — 'which vendors are connected to this incident'. No single document contains the answer. **RAG retrieval won't chain the facts for you. You need a knowledge graph.**"

- Anthropic, [March 2026](#)

 OpenAI

"This method lets your system answer **complex, multi-faceted questions** that require reasoning over several linked facts, going **well beyond what single-hop retrieval can accomplish.**"

- OpenAI, [July 2025](#)

Reach



2X LLM ANSWER ACCURACY

Build more intelligent agents with GraphRAG

New study from the **National Innovation Centre for Data (NICD)** finds GraphRAG has 2x better precision & recall.



Download the
NICD report

Combining LLMs with a Graph Database can Significantly Reduce Hallucinations and Omissions

CHRISTOPHER J. WEDGE, JOSHUA STUTTER, DANNY DIXON, JACEK CALA, National Innovation Centre for Data, Newcastle University

EXECUTIVE SUMMARY

Work by the UK's National Innovation Centre for Data has shown that integration with a graph database can significantly reduce the two most significant problems holding back the exploitation of Large Language Models (LLMs) for real-world applications: hallucinations and omissions. Combining a GPT 5.4 LLM with a Neo4j graph database achieved the highest level of truthfulness when answering 510 complex questions from a standard benchmark dataset.¹ This approach significantly exceeded conventional Vector RAG and Zero-Shot LLM approaches when compared against this measure. Importantly, the knowledge graph used was built automatically from the source data. This makes the approach relevant to real-world applications where sophisticated, manual knowledge-graph engineering is impractical.

Large language models (LLMs) have fundamentally transformed the field of Natural Language Processing (NLP), demonstrating remarkable capabilities in language understanding, reasoning, summarisation, and question answering. Their ability to generate coherent and contextually relevant responses has led to widespread adoption in many domains. However, despite these advances, LLMs and LLM-based systems remain prone to a variety of failure modes of which hallucinations and omissions are the two most concerning. The former manifest themselves when the generated content is not grounded in facts, which for many applications is unacceptable. The latter may lead to missing facts and knowledge gaps in the generated answers, which can be equally harmful. In both cases the LLM response will often appear coherent and plausible such that the user is unaware of the erroneous content. Although state-of-the-art LLMs are continuously improving in providing grounded and complete answers, they still often fail for complex questions such as those collated within the MoNaCo benchmark we tested against.

To address this challenge, we focused on one component of a typical LLM-based system – retrieval augmented generation (RAG). RAG is a very effective mechanism to contextualise the question at hand and assist LLMs to generate answers that are complete and grounded in facts. We explored a simple method to improve the quality of answers generated by LLMs. A typical vector-based RAG system was augmented with a lightweight graph database extracted from documents stored in a knowledge base. Importantly, the graph was not a rich, manually curated semantic knowledge graph. Instead, it was built from article titles, section titles, paragraphs, chunks, redirects, and links between articles or paragraphs. This makes this approach relevant to real-world applications where sophisticated knowledge-graph engineering may be impractical.

To investigate whether queries over structured knowledge bases can improve the performance of complex question answering when compared to a conventional, unstructured RAG, we set up a pipeline that ingested the English Wikipedia dataset into a Neo4j knowledge graph, storing both vector embeddings needed for semantic search and additional document structure information. The system was then evaluated on 510 complex questions from the MoNaCo

¹T. Wallson, et al. MoNaCo: More Natural and Complex Questions for Reasoning Across Dozens of Documents. *Transactions of the Association for Computational Linguistics* 14 (Jan. 2020), 23–46. <https://doi.org/10.1162/TACL.a.64>



I need my agents to **act reliably at scale**
across my landscape of **disparate data sources**

1

Meaning

2

Reach

3

Learning

**Both of these need
to be self learning**






Because the enterprise landscape
is constantly changing

Because the enterprise landscape
is constantly changing


And because your agents need
to wake up smarter tomorrow
than today

Three memory types, one knowledge graph

Short-term, long-term, and reasoning memory are connected nodes in the same Neo4j graph — not three separate stores




From conversation to knowledge to action




Skill distillation: from memory graph to grounded Agent Skill

Turn a scoped slice of a workspace's memory graph into a portable, provenance-grounded Agent Skill — procedure as a typed execution graph, not prose



Three ideas carry the whole design



Read the blog
"From Agent
Memory to
Portable
Skills"



I need my agents to **act reliably at scale**
across my landscape of **disparate data sources**

1

Meaning

ONTOLOGIES

2

Reach

MULTI-HOP

3

Learning

MEMORY



F25 bank




Global tech platform




Leading fintech


The knowledge layer for your enterprise AI!



Ontology-based Semantic Layer



... adding memory







THE Missing Layer

Watch on  YouTube

... or adding data



Example of an agent empowered by a knowledge layer



ANNOUNCING

The Neo4j Context Engine



Business context becomes an owned enterprise asset, the way data itself did.

Today

Every agent re-derives what exists, what it means, where it lives, and what it may act on.

With Context Engine

That understanding is built once, governed, and reused by every agent and every team.

Future

Enterprises manage business meaning the way they manage data platforms now, with owners, versions, and audit.

CONTEXT ENGINE PLATFORM OVERVIEW



Studio

Control Plane

Build and manage the Enterprise Knowledge Layer.

AI & Human Validated

AI-assisted mapping proposals with human validation workflow.

Environment Control

Manage ontology, agents, execution tools, and enterprise environments.

Business Ontology

Customer Proven

Proven ontology derived from 100s of customer engagements.

Enterprise Connected

Connects business process, semantics & physical systems seamlessly.

Shared Context

Provides unified business context for AI agents and applications.

Build & Runtime

Build Layer

- Operate Knowledge Layer
- Connect catalogs & systems
- Governed & change-aware

Runtime Layer

- Headless API & CLI access
- Enterprise GraphRAG
- Serve trusted agent context

Memory

Enterprise Memory

Captures agent decisions and historical experiences into memory.

Distill Knowledge

Distills repeated experiences into enterprise-reusable knowledge.

Continuous Learning

Feeds learnings back to continually improve future agent decisions.

PLATFORM ARCHITECTURE



A six-layer meta-model that stays still

L1

Business

Process · Activity · Task · Decision · Rule · Organization

L2

Domain semantics

Concept · Attribute · Relationship · Taxonomy

L3

Data products

DataProduct · Dataset · Metric · Contract

L4

Physical

System · Database · Table · Column · API · File

L5


Capabilities

Application · Tool · API · Function

L6

Runtime


Agent · Execution · ToolCall · Outcome · Feedback · Experience



SOUNDS INTERESTING?

We're looking for
Design Partners





AI is not constrained by model intelligence

Your data architecture is the world your agents think in

Graph is the perfect data architecture for AI

Thank you