neo4j

Connecting the dots with context graphs

Stephen Chin (@steveonjava)
VP of Developer Relations @ Neo4j

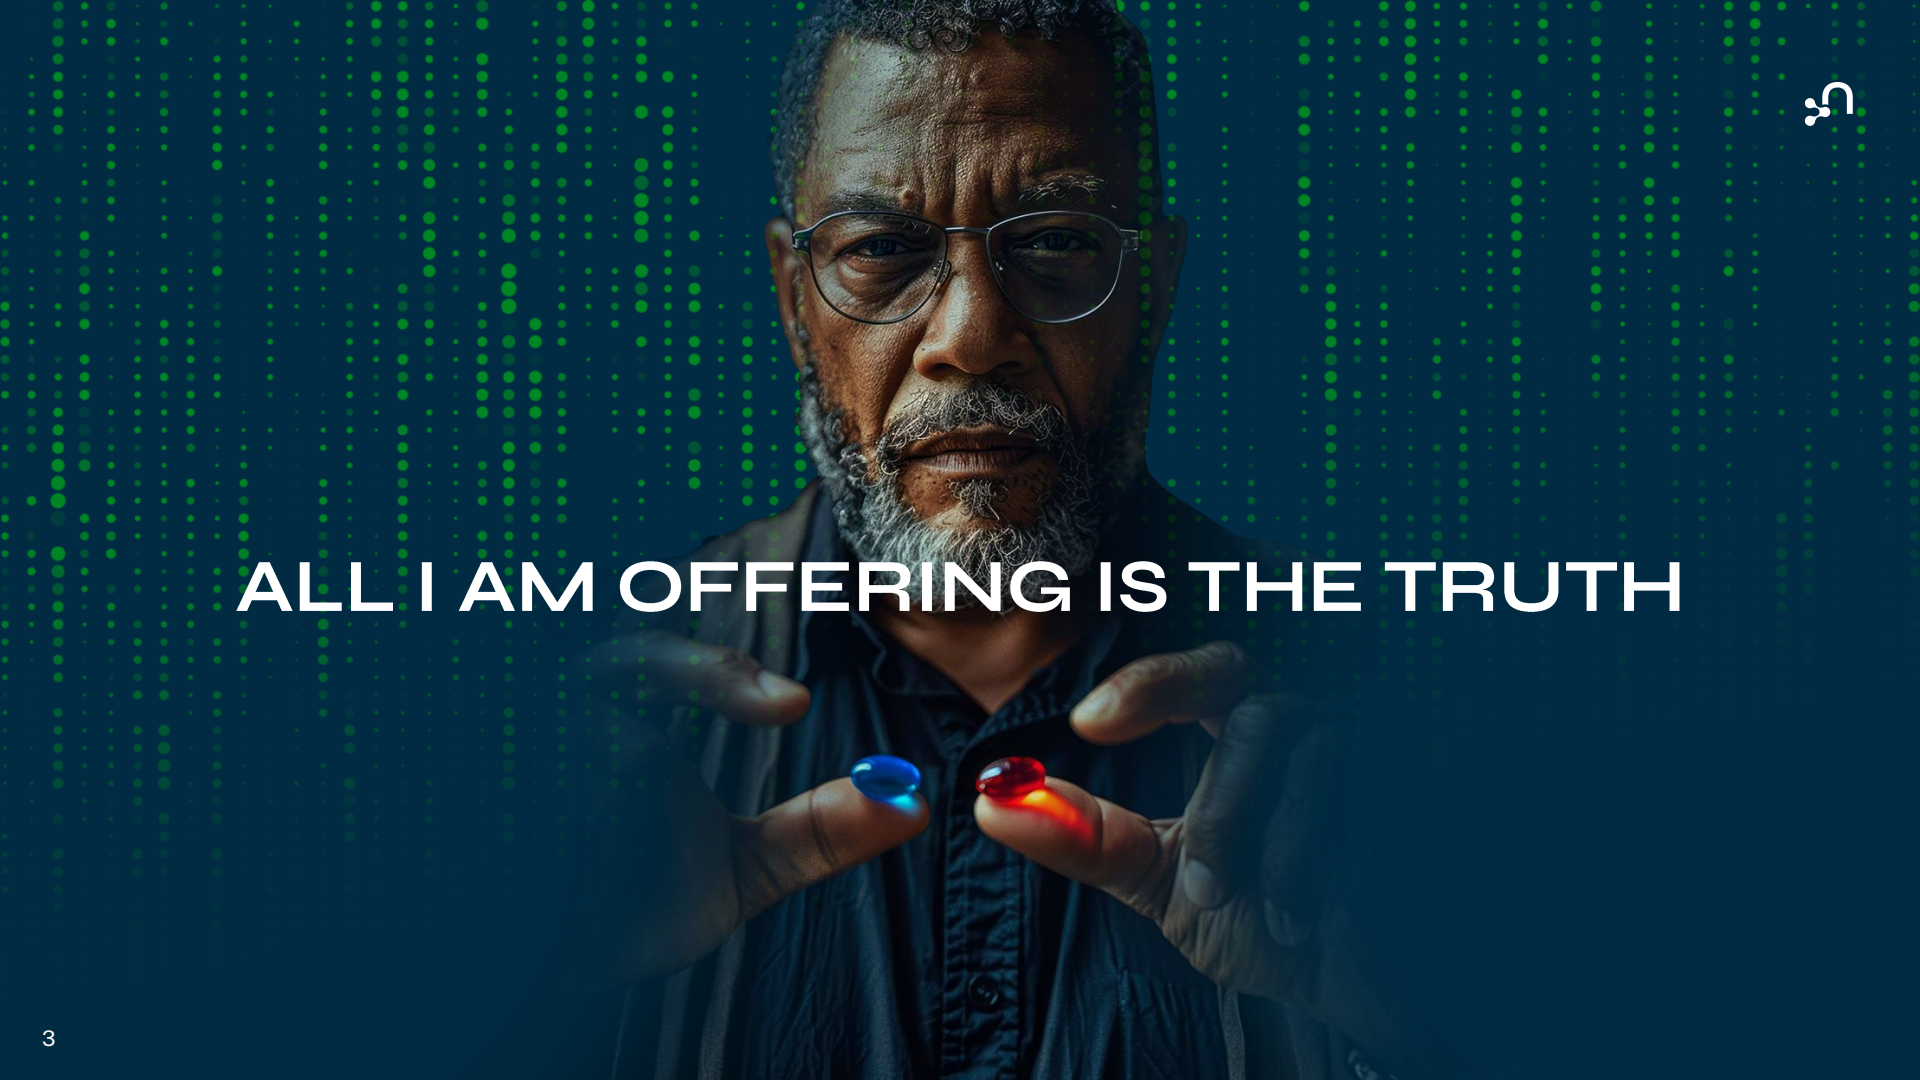
ALL I AM OFFERING IS THE TRUTH

Today

Scattered, lost reasoning

Context Lose in People's Heads & Silos
Tribal Knowledge, Cross-System Synthesis, Approval Chains Outside Systems.

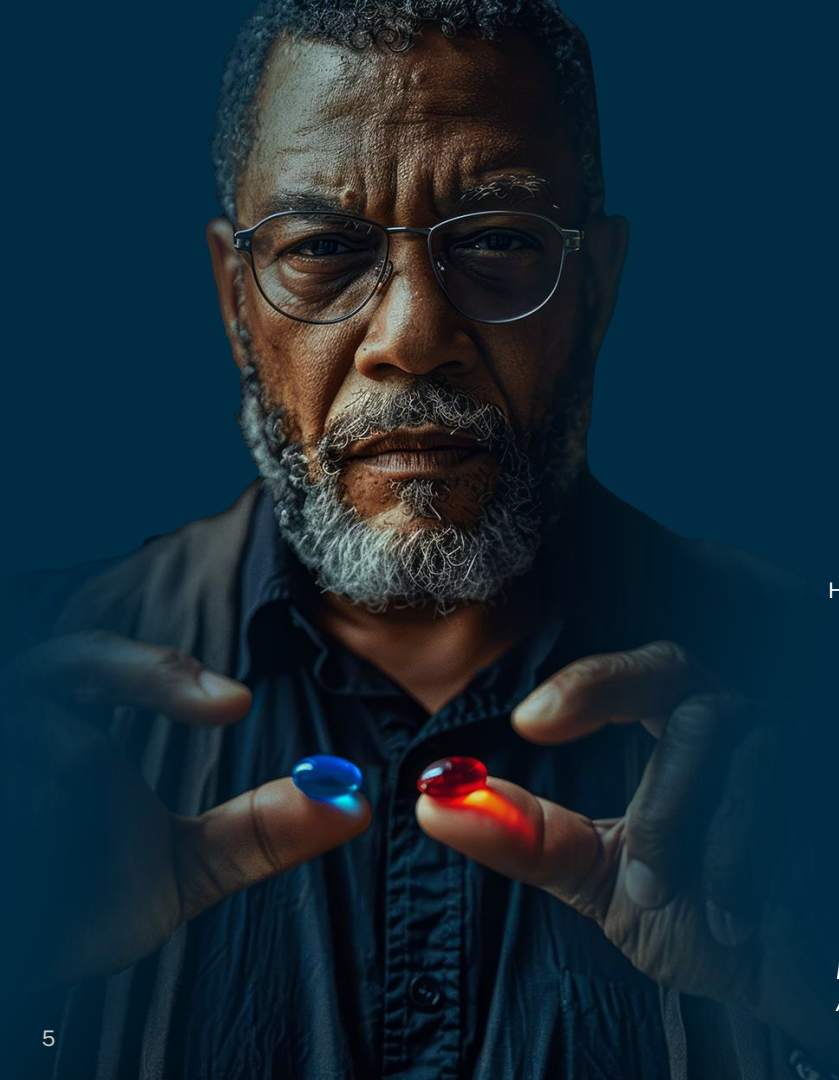

Context Graph

Connected, captured decision traces

Inputs, Policies, Exceptions, Precedents, Approvers, Approvers, Outcomes.

Decision Trace as First-Class Data

Hype cycle for Agentic AI

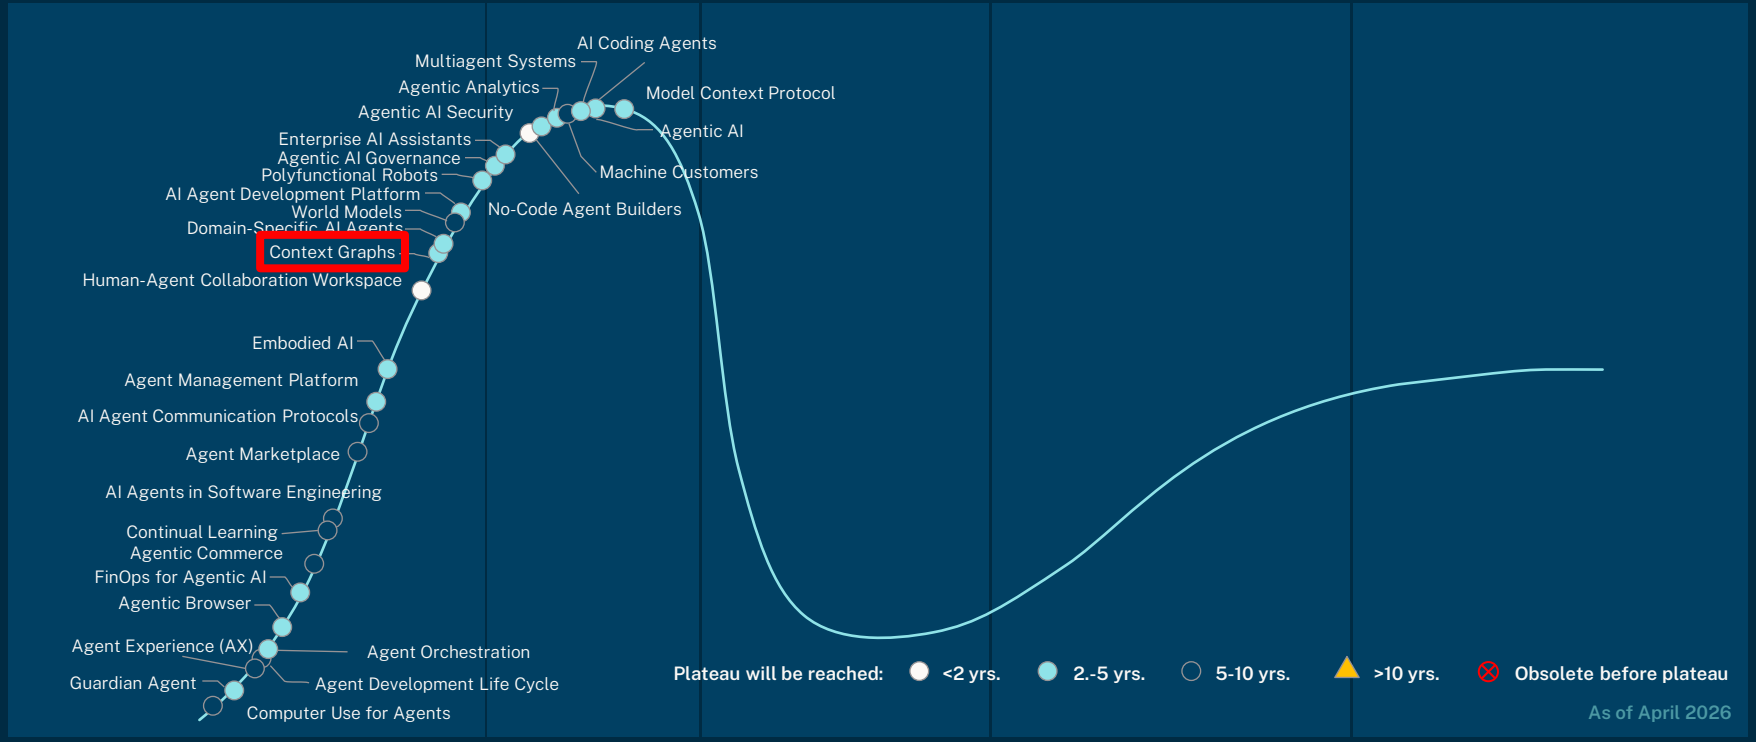

Gartner Hype Cycle for Agentic AI Report ID G00842058 – April 2026

Knowledge Graphs

Knowledge graph components

Nodes

represent entities in the graph

Relationships

represent associations or interactions between nodes

Properties

represent attributes of nodes or relationships including vectors, can be indexed

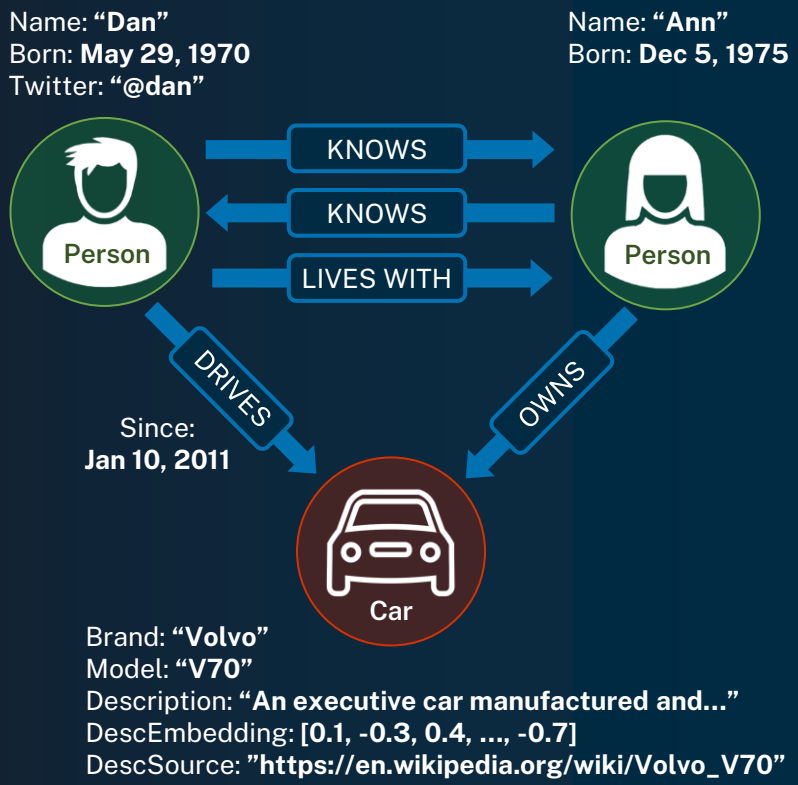

Context from graphs enables explainable AI

How do you ensure a high-quality production environment with Agents?

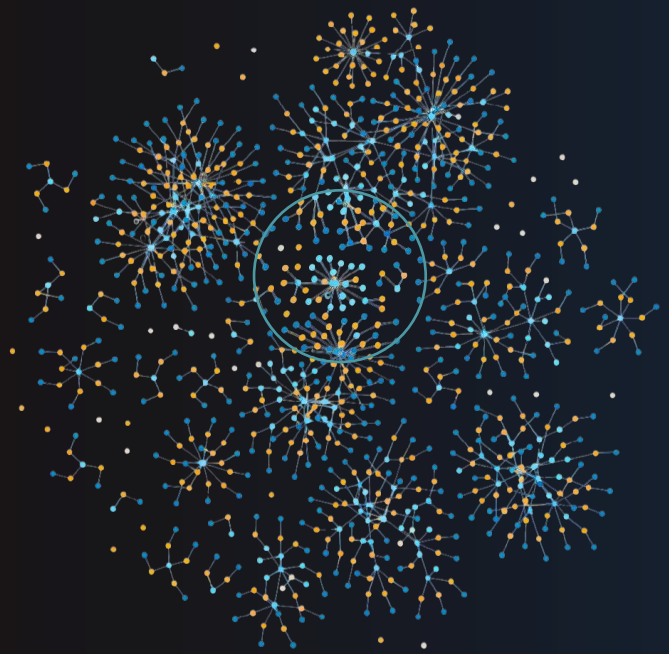

Storing

learnings from user/agent interactions in the graph database as context

Visualizing

conversations, flows, reasoning within context

Analysing

Context data for agent system performance, semi-automatically identify & apply opportunities for improvement

LLMs and Knowledge Graphs

Knowledge
Context
Enrichment

KGs

LLMs

Language
Reasoning
Creativity

BETTER TOGETHER

Enhance relevance with domain context

What was in the Care Plans associated with Andrea Jenkins's emphysema?

LLM Direct

OpenAI function call

The care plan associated with Andrea7 Jenkins714 emphysema likely included the following components:

Preventing further damage to the lungs: This may involve advising the patient to...

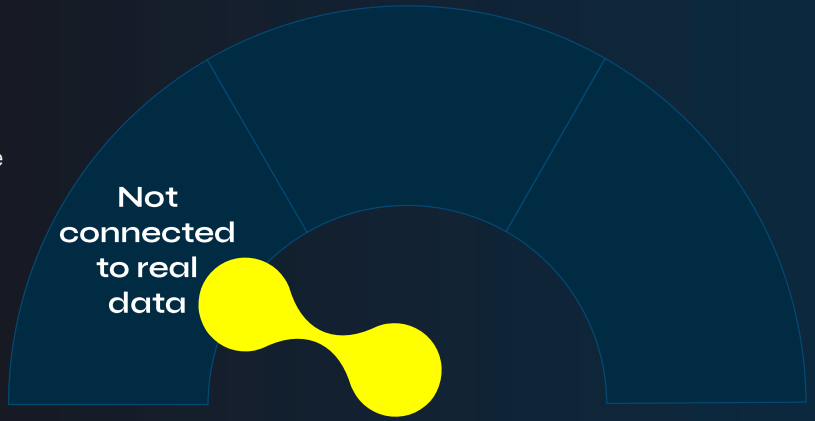

Enhance relevance with domain context

What was in the Care Plans associated with Andrea Jenkins's emphysema?

Baseline RAG

LlamaIndex-RAG backed by GPT 4.0 and Chroma vector database

The Care Plan associated with Andrea Jenkins' emphysema included activities respiratory therapy, specifically deep breathing and coughing exercises.

LLM Direct

OpenAI over Neo4j function call

The care plan associated with Andrea7 Jenkins714 emphysema likely included the following components:

Preventing further damage to the lungs: This may involve advising the patient to...

Enhance relevance with domain context

What was in the Care Plans associated with Andrea Jenkins's emphysema?

Baseline RAG

LlamaIndex-RAG backed by GPT 4.0 and Chroma vector database

The Care Plan associated with Andrea Jenkins' emphysema included activities respiratory therapy, specifically deep breathing and coughing exercises.

LLM Direct

OpenAI over Neo4j function call

The care plan associated with Andrea7 Jenkins714 emphysema likely included the following components:

Preventing further damage to the lungs: This may involve advising the patient to...

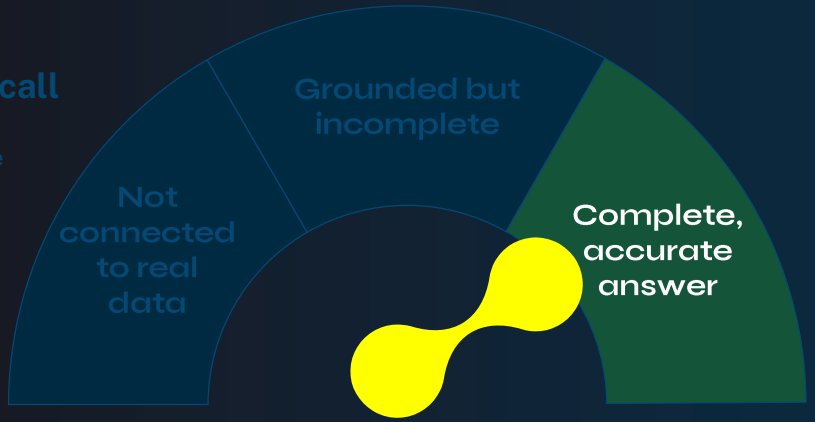

GraphRAG

Neo4j, LangChain, and OpenAI

The Care Plan for Andrea7 Jenkins714's emphysema included medication management, smoking cessation counseling, and pulmonary rehabilitation exercises.


I KNOW KUNG FU

Short term memory

- Current Context
- Compression, Relevancy
- Integrate Tool Results

Long term memory

- Episodic
- Semantic / Structural
- Procedural / Instructional

Reasoning memory

- Tool Call Traces
- Reasoning Traces
- Previous Decisions

Short-term memory

Conversation history and session state with automatic entity extraction

Conversation storage

Sessions and messages persisted as graph nodes with metadata

Multi-stage entity extraction

Pipeline combining spaCy, GLiNER, and LLM extractors with configurable merge strategies

Entity resolution

Multi-strategy dedup: exact, fuzzy, and semantic matching with type-aware resolution

```
session = await memory.add_session("user_123")
await memory.add_message(session.id, role="user",
content="Review Jessica's account")
```

Neo4j agent memory entity extraction pipeline

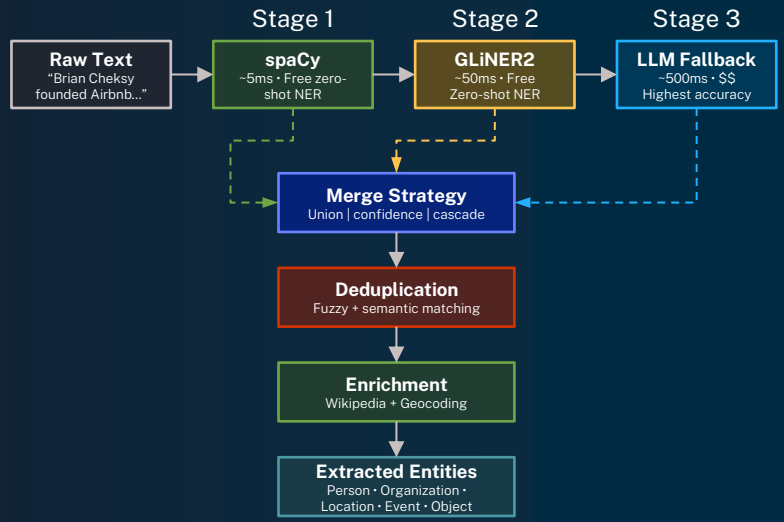

Long-term memory

Persistent knowledge graph of entities, relationships, and learned preferences

POLE+O Entity Model

Person	Customers, employees, stakeholders
Organization	Companies, teams, departments
Location	Addresses, regions, countries
Event	Transactions, decisions, meetings
Object	Accounts, products, documents

Key Capabilities

- Cross-conversation knowledge persistence
- Temporal relationships (valid from/to)
- Configurable entity subtypes
- Automatic extraction from conversations
- Vector search across all entity types

Reasoning memory

Decision traces, tool usage audits, and provenance — the layer that makes AI explainable

Tool call traces

Every tool invocation, parameters, results — complete audit trail

Decision provenance

Why did the agent choose this path? Causal chain fully recorded

Learning from experience

Agent checks if it solved something similar before, reuses successful patterns

Compliance & debugging

When something goes wrong, you can trace exactly what happened and why

OPERATOR, WE NEED AN EXIT

Why graphs for agent memory?

Relationships Are First-Class

Connections between entities, decisions, and events are the data — not an afterthought

Multi-Hop Traversal

Follow chains: Customer → Account → Transaction → Decision → Policy → Employee

Structural Similarity

Graph embeddings (FastRP) find similar situations by network topology, not just text

Explainable Decisions

Full provenance chain: trace exactly how and why every decision was made

Cross-Session Knowledge

Information learned in one conversation is available in all future interactions

Production Ready

Neo4j: ACID compliance, enterprise scale, proven technology

Neo4j-agent memory

Complete Agent Memory API

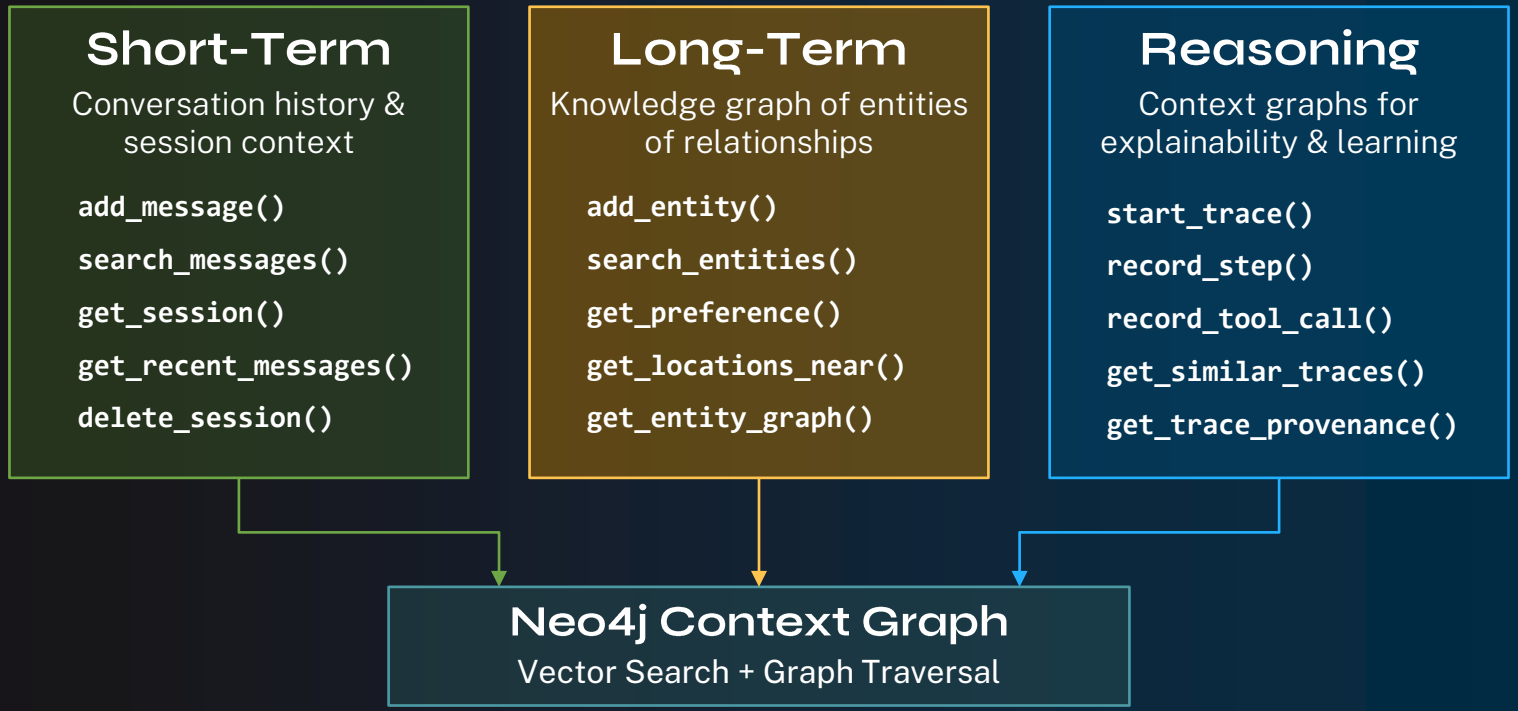

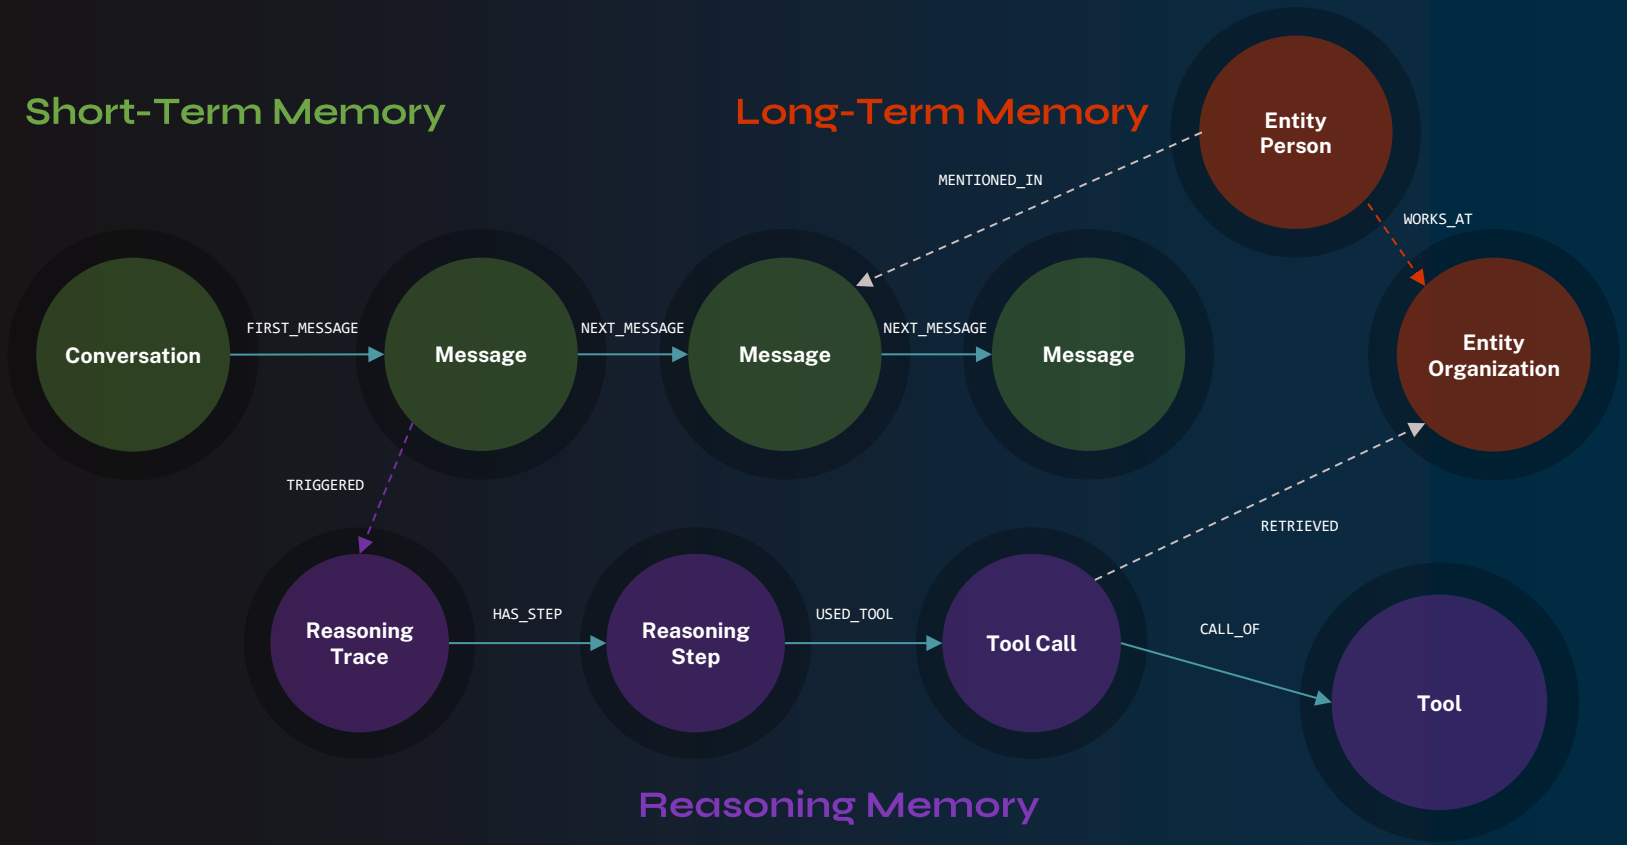
Short-Term Memory
Long-Term Memory
Entity
Person
MENTIONED_IN
WORKS_AT
Conversation
FIRST_MESSAGE
Message
NEXT_MESSAGE
Message
NEXT_MESSAGE
Message
Entity
Organization
TRIGGERED
RETRIEVED
Reasoning
Trace
HAS_STEP
Reasoning
Step
USED_TOOL
Tool Call
CALL_OF
Tool
Reasoning Memory

Lenny's memory architecture

What the Agent Does

Knowledge Graph Search

Traverse entities (people, companies, topics) and their relationships across episodes

Conversation Memory

Remembers your history and learned preferences, tailors recommendations

Reasoning Reuse

Checks if it solved something similar before, reuses successful patterns

Interactive Visualization

Explore how guests, companies, and topics connect — tool call result cards

Lenny’s memory demo

Conversations

+ New Conversation

Quick Start

Click a topic to start a new conversation

what are the Three Memory Types Every Production AI Age...
4/8/2026

Use memory graph search to explore what guests say about...
4/8/2026

Show me locations mentioned in the Brian Chesky episode
4/8/2026

Use memory graph search to

Powered by

@neo4j-labs/agent-memory

Neo4j | Lenny's Podcast | Blog | Source

Lenny's Memory Beta

About | GitHub

Ask about Lenny's Podcast

Explore insights from 299 podcast episodes with AI-powered graph memory. Click a topic or type your own question.

Explore a Topic
Use memory graph search to explore what guests say about scaling engineerin...

Who is Brian...
Tell me about Brian Chesky

Top Companies
What are the most mentioned companies in the podcast?

Related Entities
What entities are related to Airbnb?

Speaker Quotes
What did Brian Chesky say about micromanagement and being in the details?

Map View
Show me locations mentioned in the Brian Chesky episode

Ask about Lenny's Podcast...

Press Enter to send, Shift+Enter for new line

Agent Configuration

Agent Capabilities

Multi-step Reasoning
Plans and executes complex queries step by step

Conversation Memory
Maintains context across messages in the thread

Preference Learning
Adapts responses based on your stored preferences

Knowledge Graph
Queries entities and relationships in Neo4j

Available Tools 21

Tool Call Cards 7


WE NEED GRAPHS . . . LOTS OF GRAPHS

What is a context graph?

A knowledge graph specifically designed to capture decision traces — the full context, reasoning, and causal relationships behind every significant decision.

Traditional Audit Log

Records actions only

"Transaction rejected at 14:32"

No relationships between events

No causal chain or reasoning

Flat, disconnected records

Context Graph

Captures the full "why"

Decision traces & causal chains

Entities, relationships, events

Tribal knowledge made queryable

Connected, traversable structure

Context graph based retrieval

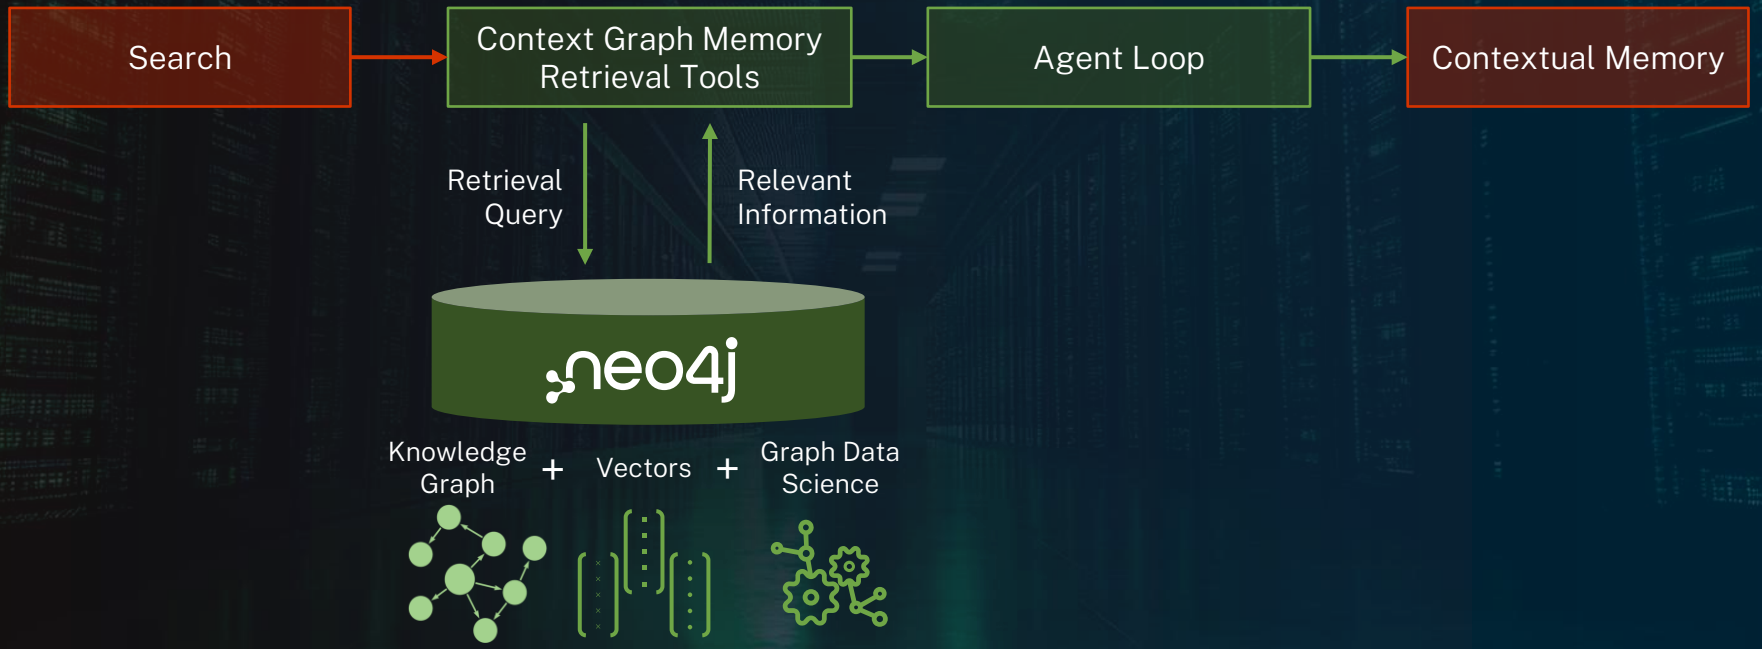

Context graph: financial services data model

ENTITIES (What Exists)
Person, Account, Transaction, Organization

EVENTS (What Happened)
Decisions, Transactions, Approvals, Rejections

CONTEXT (The Why)
Policies Applied, Risk Factors, Employee Reasoning

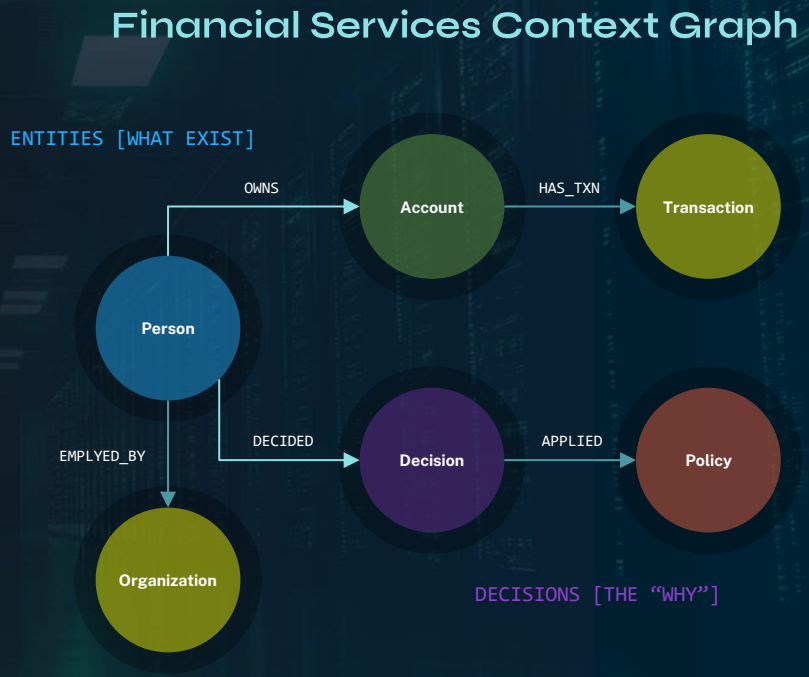

Blog: neo4j.com/blog/agentic-ai/hands-on-with-context-graphs-and-neo4j/

Context graph demo architecture

Ai-powered decision tracing for financial institutions

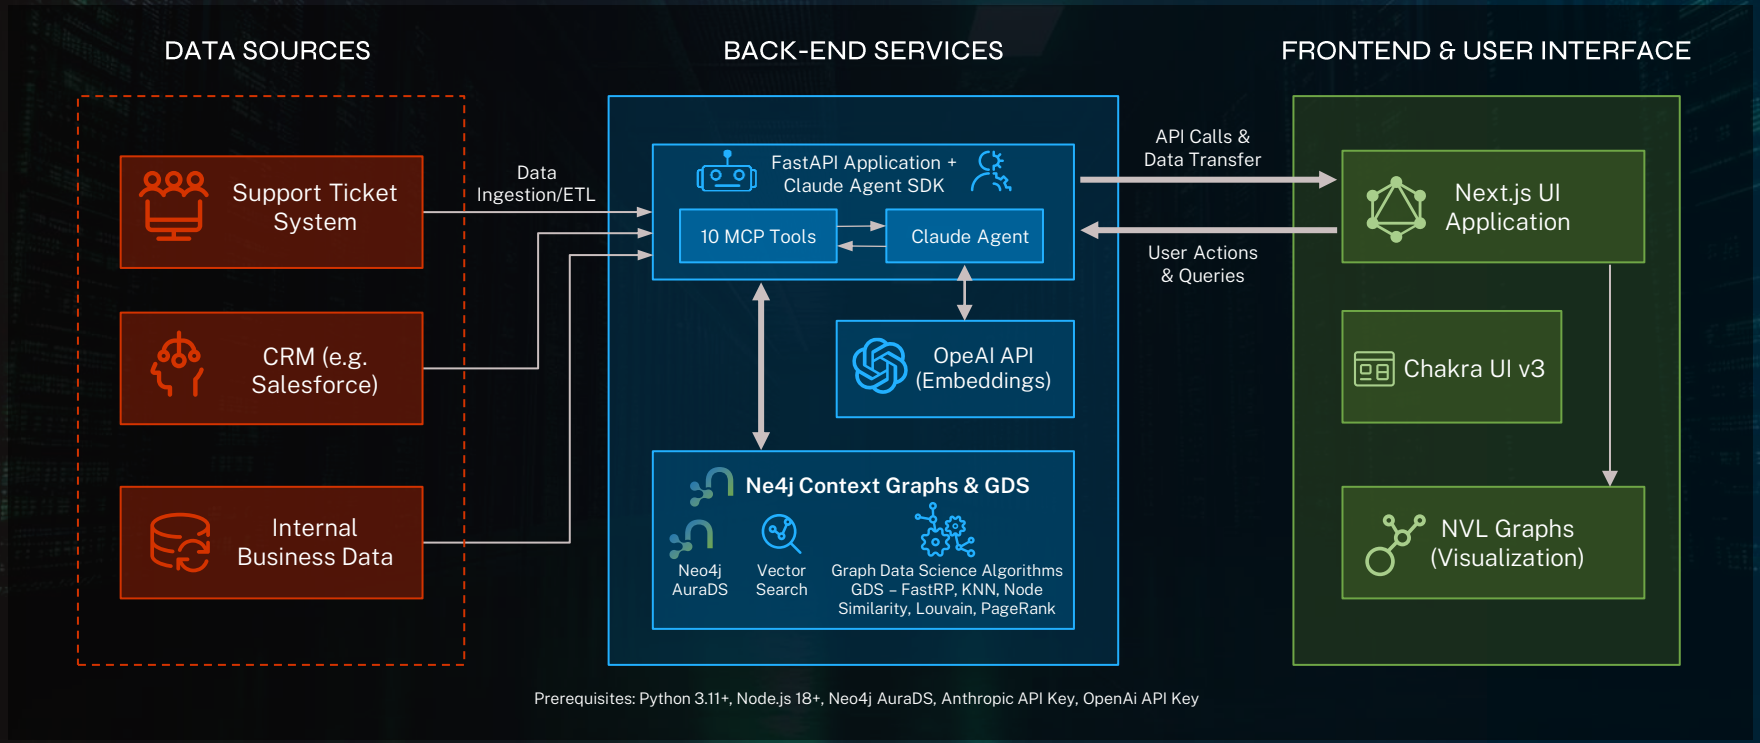

Context graph demo

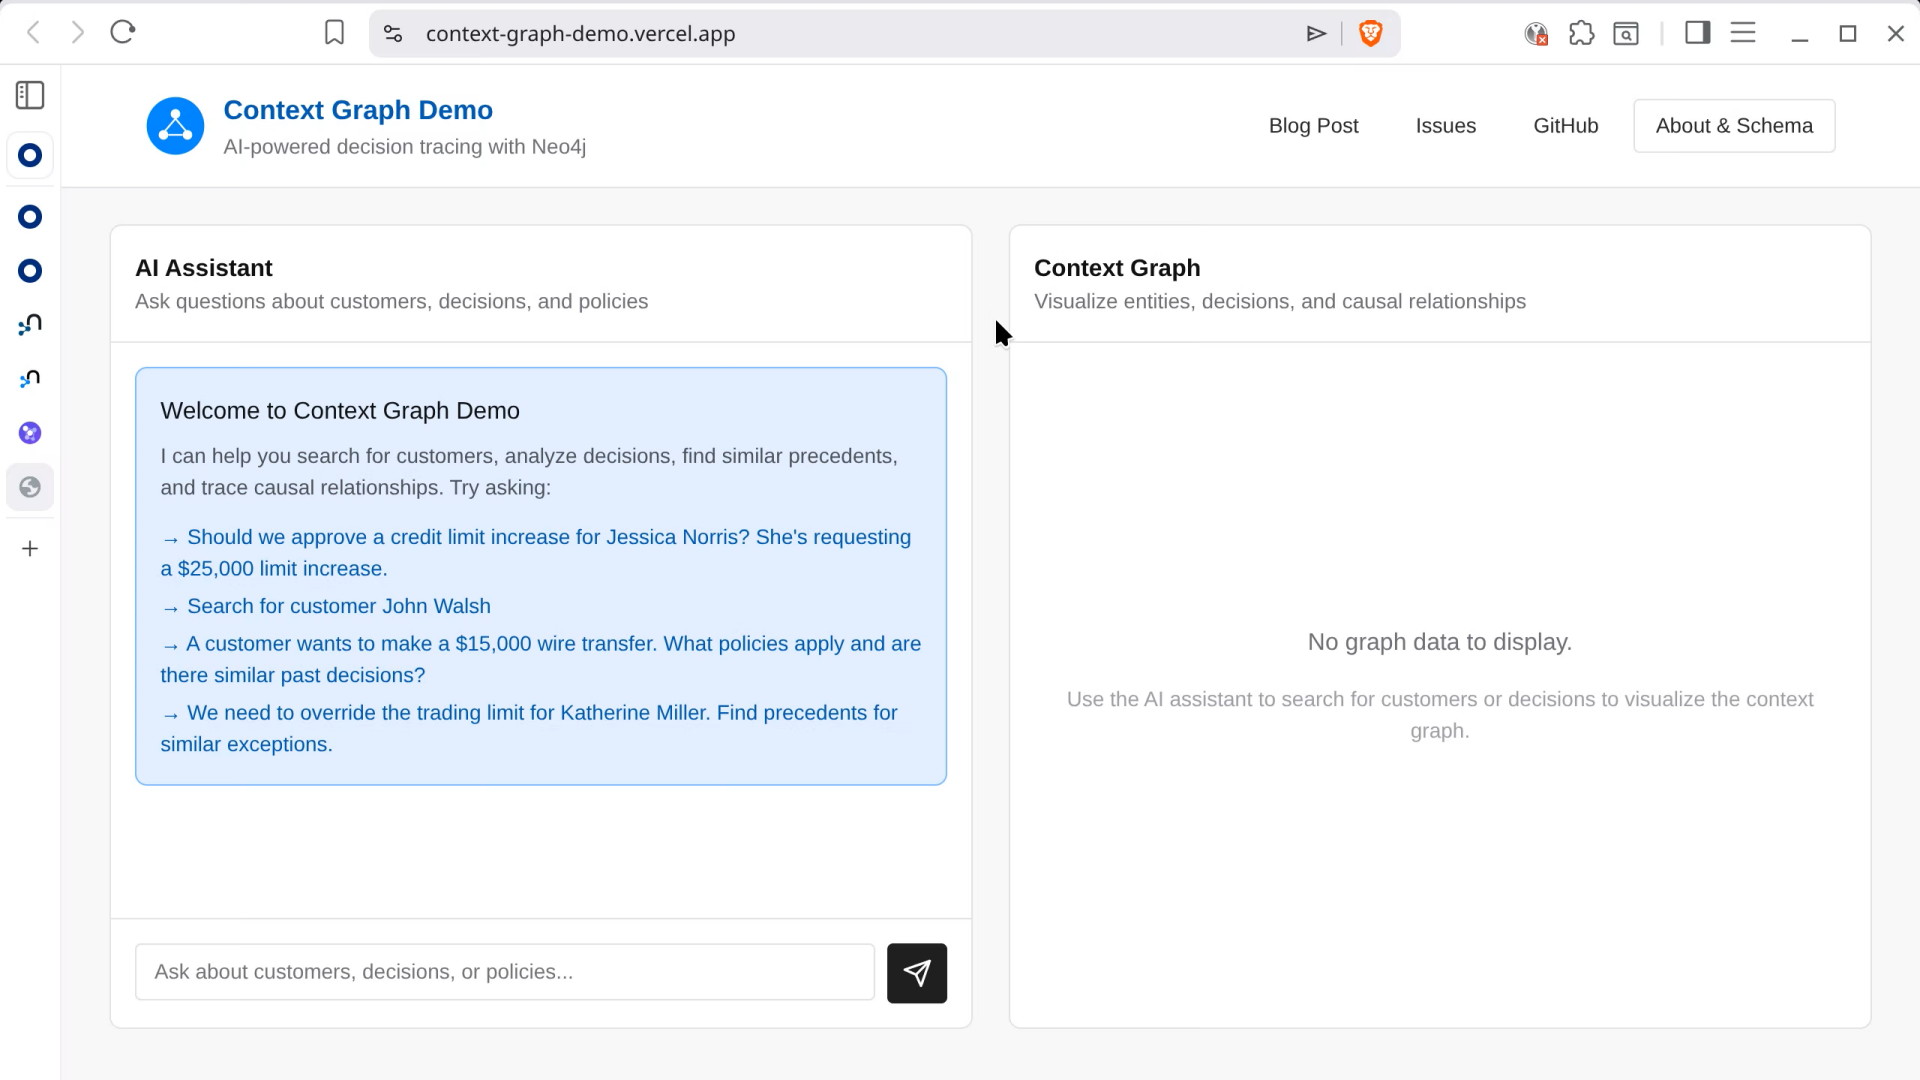
context-graph-demo.vercel.app
Context Graph Demo
AI-powered decision tracing with Neo4j
Blog Post
Issues
GitHub
About & Schema
AI Assistant
Ask questions about customers, decisions, and policies
Welcome to Context Graph Demo
I can help you search for customers, analyze decisions, find similar precedents, and trace causal relationships. Try asking:
→ Should we approve a credit limit increase for Jessica Norris? She's requesting a $25,000 limit increase.
→ Search for customer John Walsh
→ A customer wants to make a $15,000 wire transfer. What policies apply and are there similar past decisions?
→ We need to override the trading limit for Katherine Miller. Find precedents for similar exceptions.
Ask about customers, decisions, or policies...
Context Graph
Visualize entities, decisions, and causal relationships
No graph data to display.
Use the AI assistant to search for customers or decisions to visualize the context graph.

neo4j

Graph Academy

Learn everything you need to know to combine Generative AI and knowledge graphs to produce highly accurate responses, with rich context and deep explainability.

https://dev.neo4j.com/ga-rag

Tomorrow's AI Unicorns Build on Neo4j

AI is the most ambitious startup frontier in decades. To succeed in this environment, the next unicorn founders need more than ambition.
They need speed, reliability, and architecture that adapts as fast as AI evolves.
The Neo4j Startup Program equips founders with everything they need to build explainable, scalable, production-ready AI applications.

Founders get support across three core areas:

Runway: Receive credits for Neo4j Aura, a graph-native database built for AI.

Expertise: Validate your GenAI architecture with Neo4j graph engineers.

Exposure: Co-marketing and partnership opportunities to drive visibility and product adoption.

From MVP to IPO, Neo4j is your trusted AI innovation partner — every build, every launch, every win.

neo4j

Thank you

Stephen Chin (@steveonjava)
VP of Developer Relations @ Neo4j