neo4j Graph Intelligence Platform

graphsummit

Product vision & roadmap

Sudhir Hasbe
President & Chief Product Officer, Neo4j
Sudhir.hasbe@neo4j.com

Enterprise agents need a knowledge layer

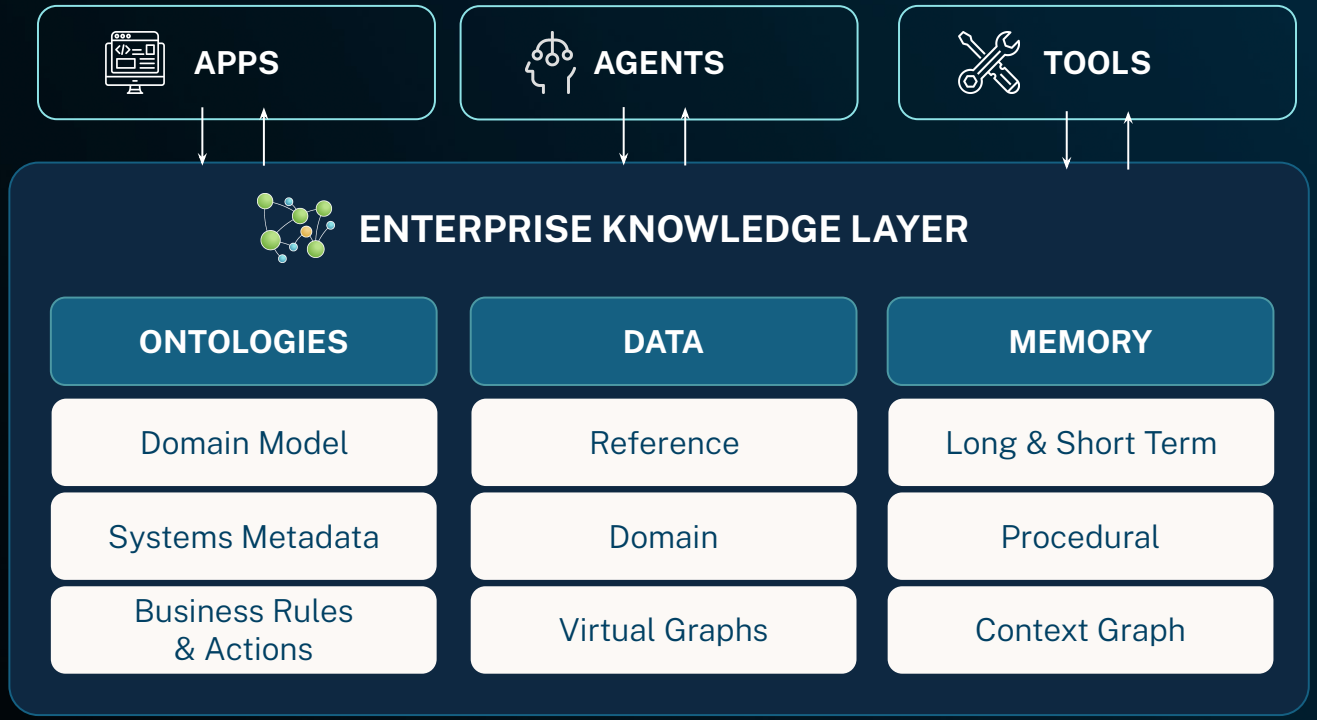

An **Enterprise Knowledge Layer** provides the semantic layer and connected context necessary for agents to move from simple pattern matching to accurate, grounded reasoning.

Benefits of graphs based knowledge layer in agentic AI

1

Multi-hop reasoning enables agents to perform complex tasks required in enterprise use cases

2

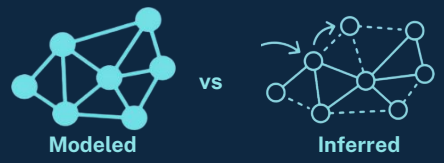

Modeled relationships let an agent reason across hops and prove the path it took

3

GraphRAG delivers more accurate results than vector only RAG

Product Vision

What Premium, trusted cloud-native **graph intelligence platform**

Why Transform data to knowledge and power Enterprise Agentic AI

How Cross cloud, scalable, easy to use and cost effective

Who For developers, data analysts, and data scientists

Neo4j’s Graph Intelligence Platform

The most comprehensive platform to build an enterprise knowledge layer

GraphAware Solution	Purpose-built domain applications delivering turnkey intelligence and knowledge discovery.
Context Engine	Neo4j Context Engine is the platform for building and operating a shared, governed Knowledge Layer for AI agents.
Graph AI	Enterprise-grade graph algorithms & agents that uncover hidden patterns and predictive insights.
AI Powered Graph Tools	Intelligent interfaces and query assistants that streamline graph development, visualization, and management.
Neo4j Databases	Industry's most comprehensive graph database portfolio for any scale.

Neo4j's Graph Intelligence Platform

The most comprehensive platform to build an enterprise knowledge layer

GraphAware Solution

- Intelligence Analysis
- Financial Crime

Context Engine

- Ontology
- Studio
- Build Services
- Runtime
- Memory

Graph AI

- Analytics (Algorithms)
- Aura Agents
  Build multi hop agents in seconds.
- Agentic Ecosystem
  Integrate: AWS, GCP, Azure, LangChain, ServiceNow, SalesForce, Databricks & Snowflake

AI Powered Graph Tools

- Query
- Bloom
- Dashboards
- Data Import
- Doc Intel
- Ontology Mgmt

Neo4j Databases

- Enterprise Edition
- InfiniGraph
- AuraDB
- Virtual DB
- Composite DB

Neo4j's Graph Intelligence Platform

The most comprehensive platform to build an enterprise knowledge layer

GraphAware Solution

- Intelligence Analytics
- Financial Crime

Context Engine

- Ontology
- Studio
- Build Services
- Runtime
- Memory

Graph AI

- Analytics (Algorithms)
- Aura Agents
  Build multi hop agents in seconds.
- Agentic Ecosystem
  Integrate: AWS, GCP, Azure, LangChain, ServiceNow, SalesForce, Databricks & Snowflake

AI Powered Graph Tools

- Query
- Bloom
- Dashboards
- Data Import
- Doc Intel
- Ontology Mgmt

Neo4j Databases

- Enterprise
  Enterprise Edition
- InfiniGraph
- AuraDB
- Virtual DB
- Composite DB

Neo4j Databases

INFINITE SCALABILITY

Graph database for your needs

Native

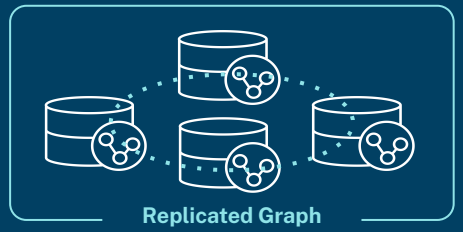

Native Graph that can be on single machine or shared across many machines to handle large scale.

Scale out for reads and large data size.

Virtual

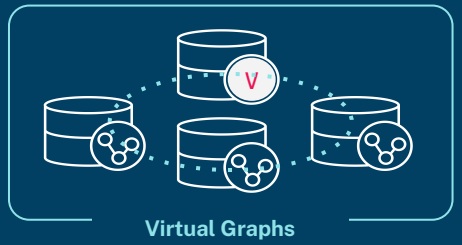

Virtual Graph that provides all graph features on top of existing data storage systems.

Support for unlimited scale but with higher latency.

Ability to materialize data.

Composite

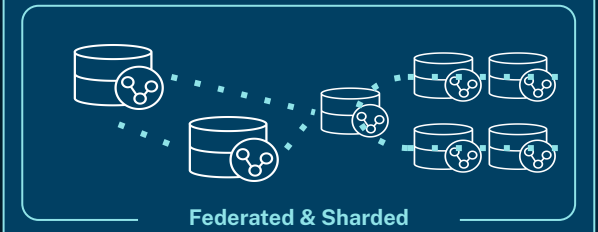

Truly unlimited scaling with any data source

Combination of sharded and federated data

Use cases: unlimited growth for mission critical, app-based queries/analysis

Neo4j Databases

INFINITE SCALABILITY

Graph database for your needs (Native Storage)

Core HA Deployment

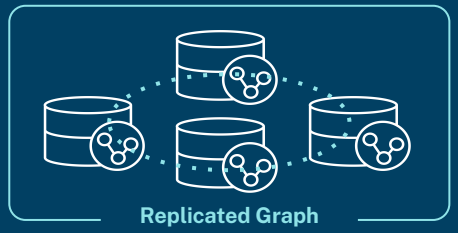

Native Graph that can be on single machine or shared across many machines to handle large scale.

Scale out for reads and large data size.

Multi-database

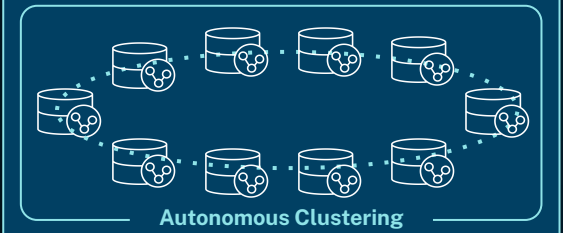

Hundreds of databases in a single cluster

Thousands of databases in multiple clusters with aliasing and user authorization / impersonation

Use cases: SaaS, internal services based on the same graph model

Infinigraph

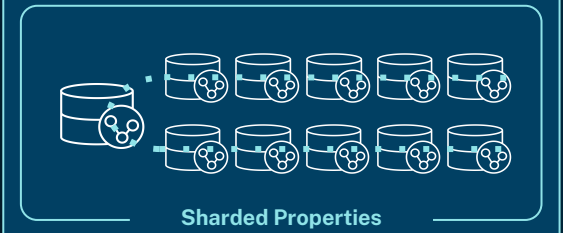

Tens of TBs of data in a graph with property sharding

Hundreds of TBs with full-text and vector data

Use cases: heavy graph analytics, user queries/analysis, seamless integration with connectors/plugins

Neo4j Enterprise Edition

CORE DATABASE ENHANCEMENTS

Neo4j Enterprise Edition

Autonomous clustering with **multi-db** for multi-tenant SaaS solutions

Block Format is superior graph native format to group graph data into blocks that reduces amount of data read with greater **performance** in memory-constrained scenarios

Parallel Runtime speeds up analytical queries up to 100x by concurrent execution on all cores

CDC enables graph changes published as events in real-time (full or diff mode)

Autonomous Clustering (with Multi DB)

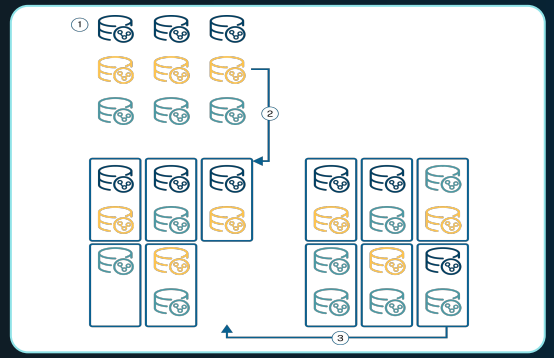

Parallel Runtime

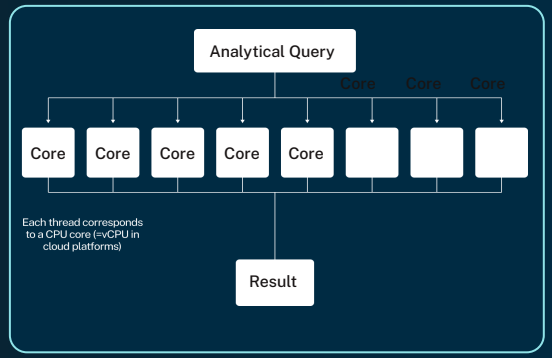

Change Data Capture

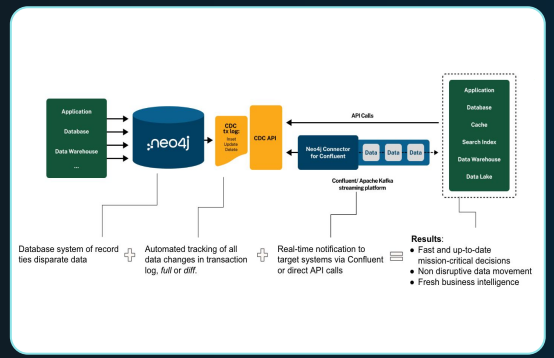

Block Format

Neo4j Databases

MOST POWERFUL DATABASE FOR GRAPHRAG

Vector Support

Vector Type in Cypher 25

Store embedding properties on nodes and relationships as vector in the database

Vector Index built into the database

Embedding generation from Cypher AI functions for OpenAI, Azure OpenAI, VertexAI, Bedrock

Vector Search with filters and Cypher SEARCH Surface (Q1 26)

IN Filters (Q2 26)
Binary Quantization (BBQ), expansion factors (Q3)

Cypher MATCH Filter, Hybrid Search (H2 2026)
Automatic Indexes (2027), Dedicated Search Instances (2027)

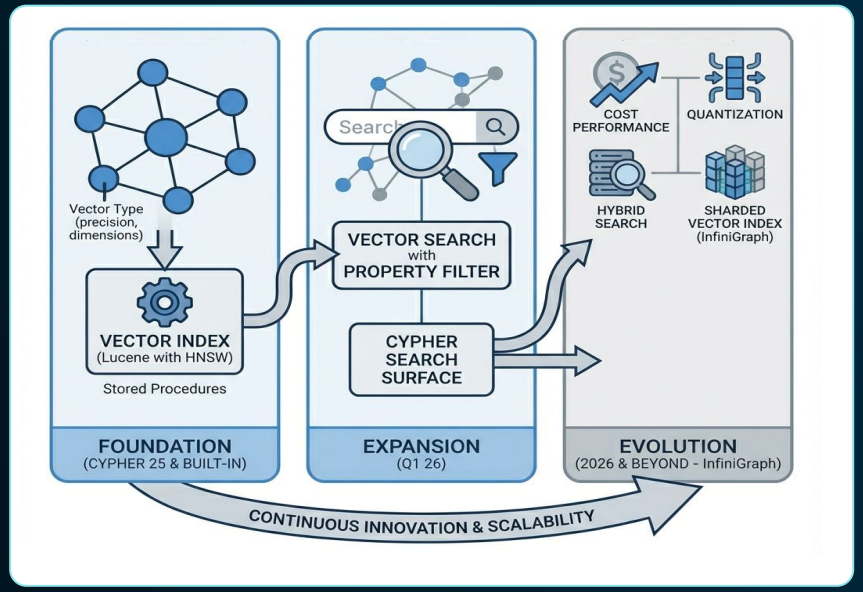

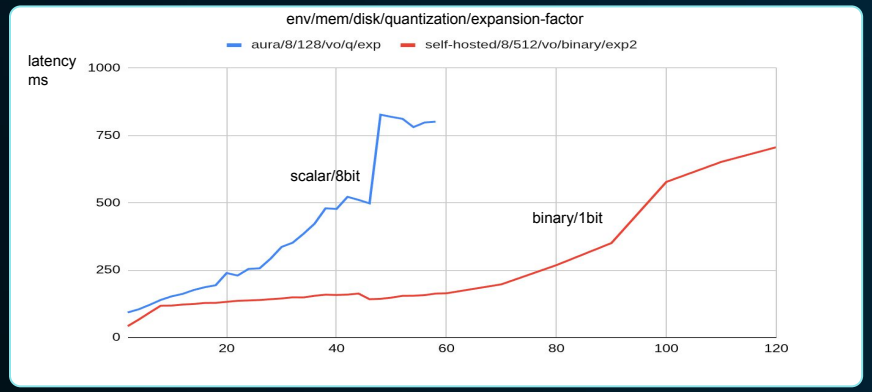

Neo4j Databases

CORE DATABASE SCALABILITY

Vector Search and Analysis

Strong improvement in the last Neo4j releases, with High Fidelity Quantized vector search

Tests with `2026.09` show better results, compared to the industry trend

Property sharding with Infinigraph proven to be the best infrastructure for vector search and graphs

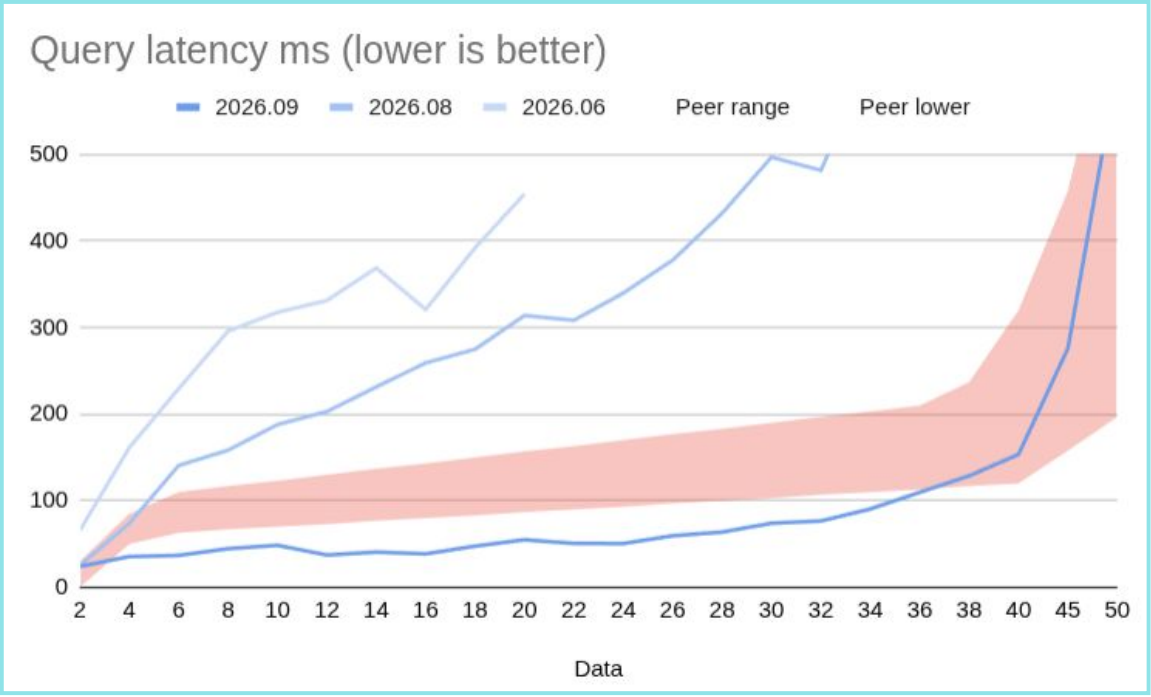

Neo4j Databases

CORE DATABASE CAPABILITIES

Cross Cluster Replication

Continuously replicate databases across separate Neo4j instances or clusters as a read-only copy.

- **Live Catch-up Protocol**
  The replicating server uses a network connection to the upstream to use the catch-up protocol
- **Backup Chain Pull**
  The replicating server continuously polls an object store bucket for diff backups

On failure of the upstream, the replica can be ‘promoted’ to a write-enabled database.

Neo4j to be Tier 1 platform in enterprises with potential to deliver RPO (0) /RTO (15min)

Core DB GA - Sept 2026

Aura - Q1 2027

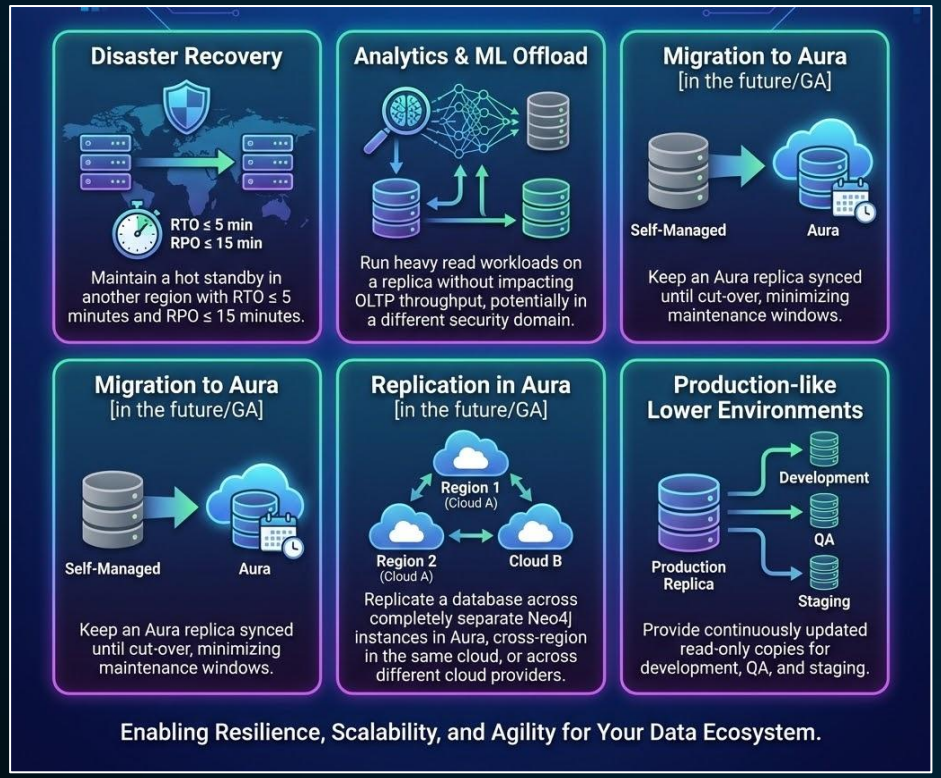

Demo
Cross Cluster Replication

InfiniGraph

Neo4j Databases

SHARDED DATABASE FOR SCALABILITY

InfiniGraph

Available

- **High data volumes** - up to 100TB (and potentially more) of graph data
- **Fast data loading** - initial bulk import and staging/live incremental updates
- **Semantic indexes** - sharded full-text and vector indexes are natively integrated
- **ACID** - fully transactional and ACID compliant
- **Simplicity** - Transparent to user, standard Cypher
- **Analytics** - supports Neo4j tools (Query, Explore, Graph Analytics)
- **Transparent** to all API calls

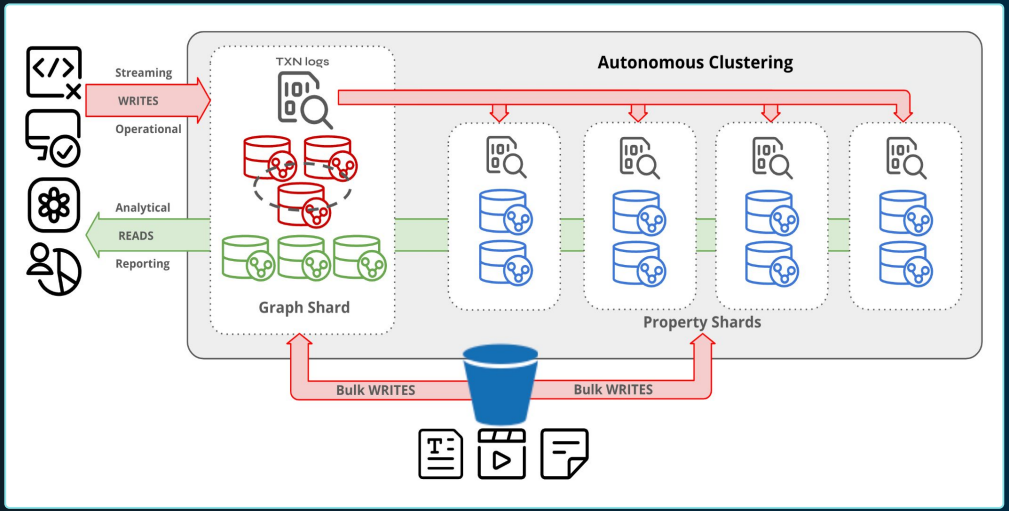

DATA INGESTION

Improved Infinigraph Distributed Import

The importer mimics the same Infinigraph architecture (in later versions it could be also used for non-sharded databases)

Use autonomous clustering to:

- Create an Infinigraph database that will be immediately available
- Write backups to object stores to load the graph on other clusters

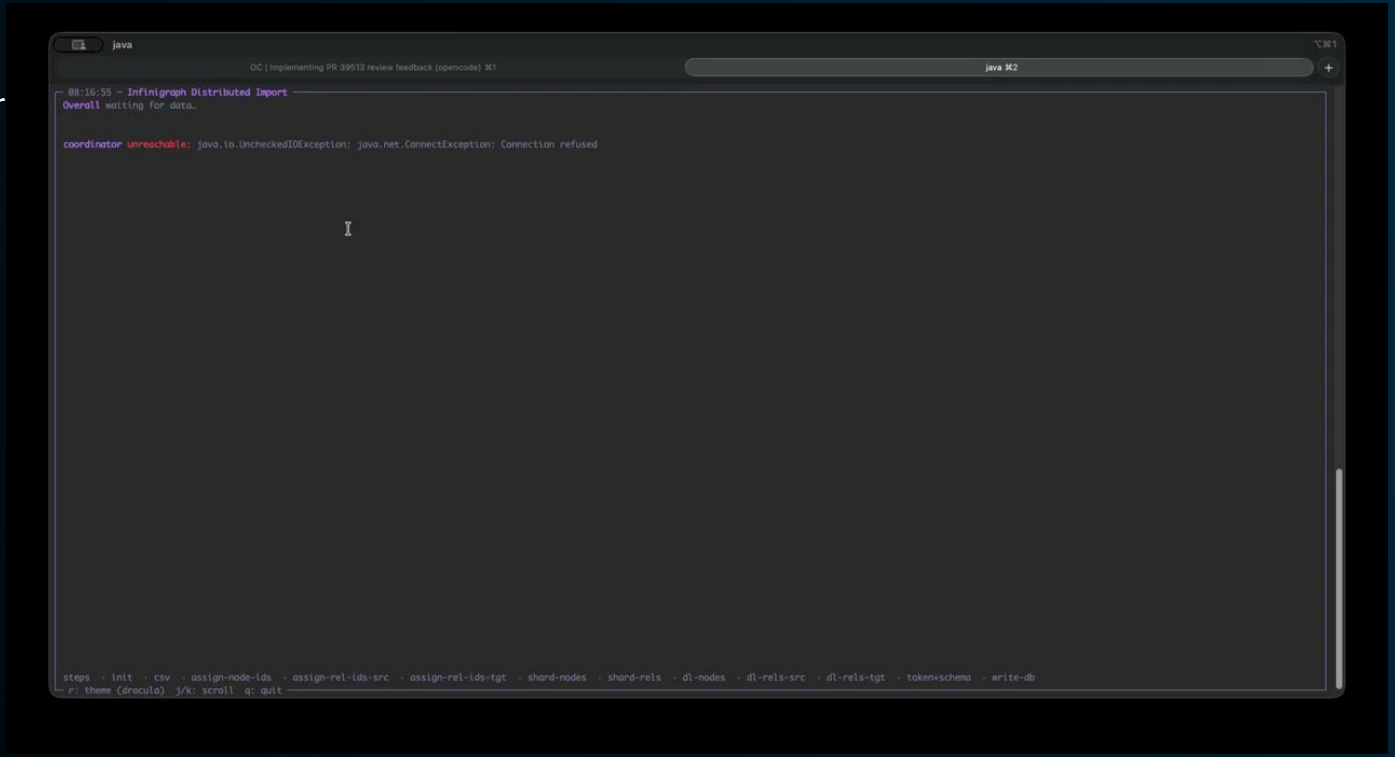

Neo4j Databases

DATA INGESTION

Improved Infinigraph Distributed Import

The importer mimics the same Infinigraph architecture (in later versions it could be also used for non-sharded databases)

Use autonomous clustering to:

- Create an Infinigraph database that will be immediately available
- Write backups to object stores to load the graph on other clusters

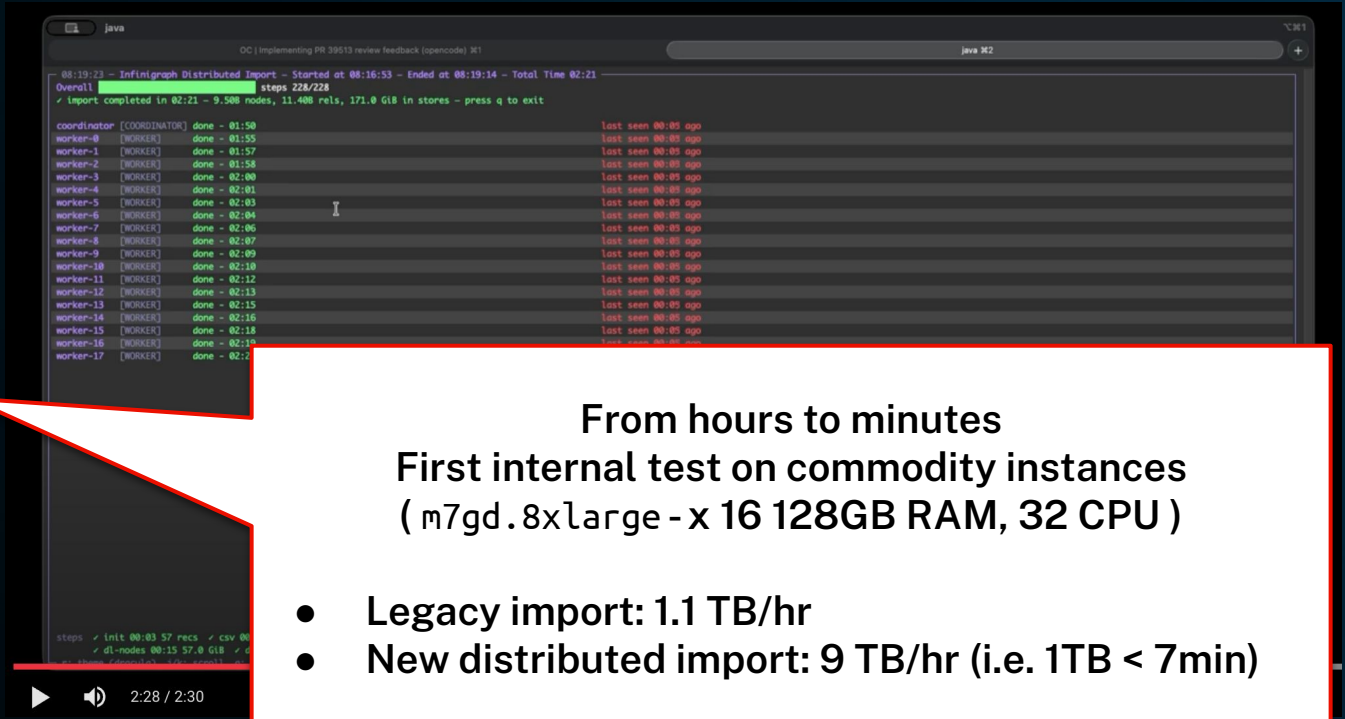

From hours to minutes
First internal test on commodity instances
(`m7gd.8xlarge` - x 16 128GB RAM, 32 CPU)

- Legacy import: 1.1 TB/hr
- New distributed import: 9 TB/hr (i.e. 1TB < 7min)

DATA INGESTION

Streaming Ingestion

`sysbench4j` Neobank test

Based on 1TB dataset/instance
`neobank_write` workload

- Graph: `m7gd.16xlarge`
  (64 vCPU ARM, 256 GiB 2TB NVMe)
- `sysbench4j`: `c7i.16xlarge`
  (64 vCPU ARM, 128 GiB)
- Optimal batch size: 100
- Optimal concurrent workers: 32
- Total of **4.3B+ Bank TXNs/day**

50k Bank TXNs/sec
4.3B+ Bank TXNs/day
for each single shard

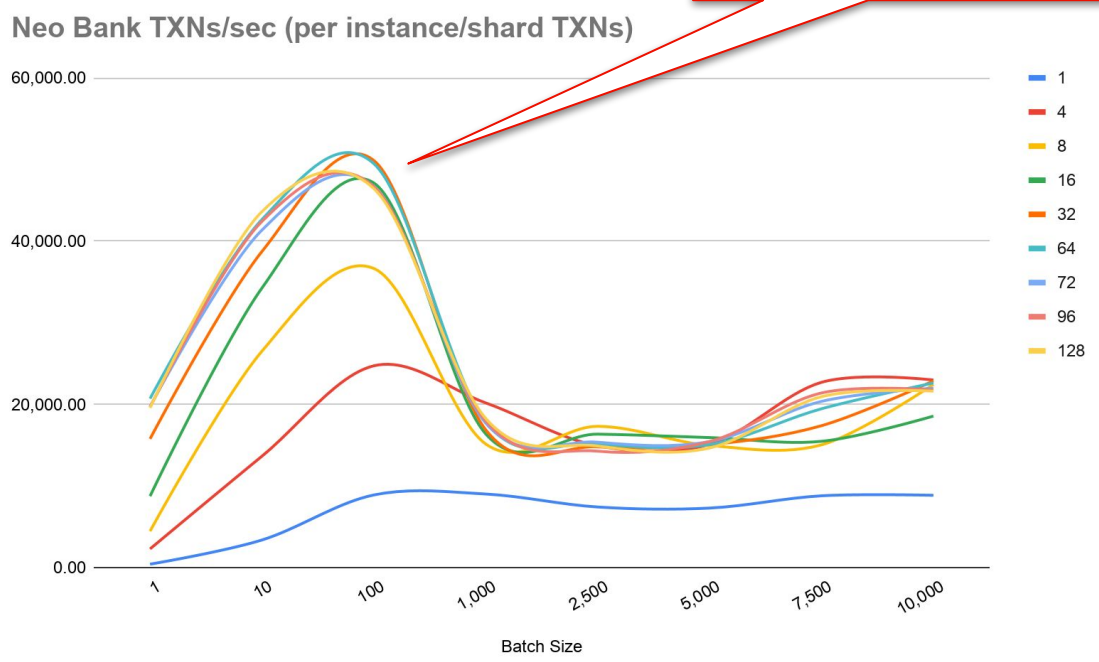

AuraDB - Cloud Managed Database

Neo4j Databases

CLOUD FIRST CAPABILITIES

AuraDB - Fully Managed Cloud DB

- Available on Google Cloud, AWS, and Azure
- 99.95% uptime SLA with self-healing cluster architecture
- Scale up for large memory-intensive graphs
- Automated backups and restore
- Automated upgrades, zero maintenance

Available on all major cloud providers

More than
30k
Managed Databases

Neo4j Databases

Comprehensive cloud offerings for all workloads

Free

Perpetually free database for customers to learn and experiment with apps

Professional

Zonal (single AZ) with functionality aimed at smaller teams & departmental solutions

Now Available

Free Trial

Business Critical

Regional (Multi-AZ) with enterprise-grade SLA & functionality

Now Available

Virtual Dedicated Cloud

Highest tier, regional (Multi-AZ) with dedicated account provisioned per customer.

Neo4j Databases

Aura Launches 2025 / 2026

Q4 '25

- New: Aura Agent (Preview)
  Date: 2 October
- New: AuraDB Pro Trial on GCP Marketplace
  Date: 6 October
- New: MultiDB (Private Preview)
  Date: 9 Oct
- New Regions: Azure Netherlands for Pro & BC
  Date: 13 October
- Feature: Aura Activity Feed
  Date: 21 October
- New: Data Importer API
  Date: 22 October
- New: Graph Intelligence for MSFT Fabric
  Date: 23 October
- Feature: Adjustable Storage - AuraDB Azure
  Date: 4 November
- Feature: Credit Balance in Aura Console
  Date: 12 November
- New AWS PAYG Marketplace listing (inc BC)
  Date: 19 November
- New: Vector Native Type in Aura
  Date: 20 November
- New: Terraform for Aura (Labs)
  Date: 3 December

Q1 '26

- New: Fleet Manager - Market Launch
  Date: 9 December
- New: Set Organization Default Cypher Version
  Date: 12 December
- New: Access backups on paused instances
  Date: 12 December
- New: Project / Billing management (Pro & BC)
  Date: 15 December
- New Regions: AWS Ohio & Sydney | BC & Pro
  Date: 17 Dec
- New: Deploy local (Desktop) databases to Aura
  Date: 17 Dec
- New Azure PAYG Marketplace listing (inc BC)
  Date: 13 Jan
- New Region: AWS Thailand (ap-southeast-7) for Virtual Dedicated Cloud
  Date: 20 Jan
- New: console.neo4j.io now the primary domain for Aura Console
  Date: 22 Jan
- New: Aura Activity Feed API
  Date: 22 Jan
- Graph Intelligence MSFT Fabric - connect existing Pro and BC running on Azure
  Date: 29 Jan
- New: Billing API (Private Preview))
  Date: 30 Jan
- New: Aura Agent GA
  Date: 4 Feb
- New: Bulk Import in Aura (Public Preview)
  Date: 24 Feb
- New: 2TB:5TB GCP (Preview)
  Date: 27 Feb
- New: Adjustable storage "1:16" all clouds / paid tiers
  Date: 2 March
- New: GRAPH TYPE in Aura (Public Preview)
  Date: 2 March
- AuraDB Self Serve Activation Codes
  Date: 10 March
- New Region: GraphQL DataAPI for VDC
  Date: 2 April
- New Explore: Multiple Scenes & Perspectives all AuraDB tiers
  Date: 8 April
- New: Connect to Desktop in Aura Console
  Date: 20 April

Q2 '26

- AuraDB: Export to Lakehouse in MS Fabric
  Date 21 April
- New: Dashboard Actions for AuraDB
  Date: 23 April
- New: Project level ML Persistence in Aura Graph Analytics
  Date 1 May
- New: Inactive Console User Pruning
  Date: 5 May
- New: 2TB RAM :5TB Store GCP (GA)
  Date: 13 May
- New: Neo4j CLI (labs) - neo4j.sh
  Date: 14 May
- New: Aura Billing - Cost Explorer
  Date: 27 May
- New: 5TB Adjustable Storage AWS
  Date: 28 May
- New: Virtual Graph (Private Preview)
  Date: 28 May
- New: Document Intelligence (Preview)
  Date: 1 June
- New: Billing API GA
  Date: 2 June
- Native Projections:Aura Graph Analytics
  Date: 9 June

Q3 '26

- New: Import supports List & Vector Data types
  Date: 3 July
- New: Invites at Organization level
  Date: 3 July
- New: Aura User Mgmt APIs
  Date 3 July
- New: SSO fully serve in Aura
  Date 15 July
- New: Fleet Manager import to Aura from CE
  Date: 15 July
- New: Virtual Graph (Preview)
  Date: 21 July
- New: MCP for Aura
  Date: 22 July
- New: Query Tabs
  Date: Aug 5
- New: Attribute Based Access Control (ABAC)
  Date: 12 Aug
- New: Graph Analytics on AuraDB Free
  Date: 13 August
- New: Instance Explorer Preview
  Date: 20 August
- New: MultiDB (Public Preview)
  Date: 25 August

Neo4j Databases

AURADB NEW CAPABILITIES

Business Critical Enhancements

Coming in 2026

Private Link connectivity: AWS first (Q4 2026 - Preview), with planned rollout to GCP and Azure.

Marketplace availability: Available via AWS, GCP and Azure marketplaces.

Security transparency: Comprehensive security white paper and reference architecture available in the Trust Center.

Global footprint: Available in 30 regions across 3 cloud providers.

Neo4j Databases

AURADB CAPABILITIES

Scalability Options

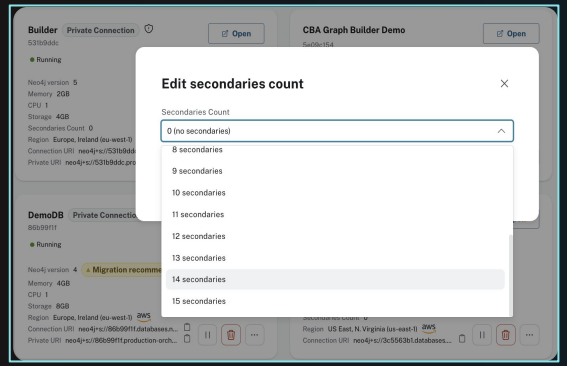

Read Scalability with up to 15 Secondaries

Now Available

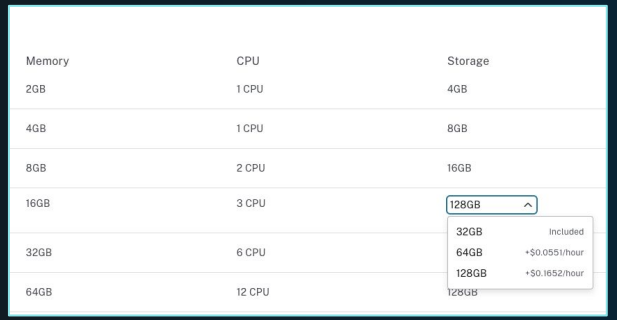

Memory	CPU	Storage
2GB	1 CPU	4GB
4GB	1 CPU	8GB
8GB	2 CPU	16GB
16GB	3 CPU	128GB
32GB	6 CPU	
64GB	12 CPU	

Storage	
32GB	Included
64GB	+$0.0551/hour
128GB	+$0.1652/hour

Adjustable Storage up to 1:16 ratio

Now Available

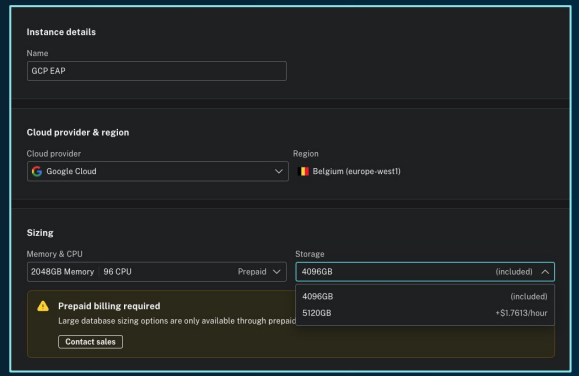

Scale Up to 2TB (RAM): 5TB (Disk) on BC/VDC

Now Available

Neo4j Databases

AURADB NEW CAPABILITIES

Multi-DB on Aura

Enables AuraDB customers to create **multiple databases** within a single AuraDB Instance.

Cost Efficiency: Reduces the need for multiple AuraDB instances, lowering infrastructure and operational costs.

Improved Scalability: Allows users to add databases quickly without provisioning new instances, speeding up growth.

Public Preview - Available
GA Coming Q4

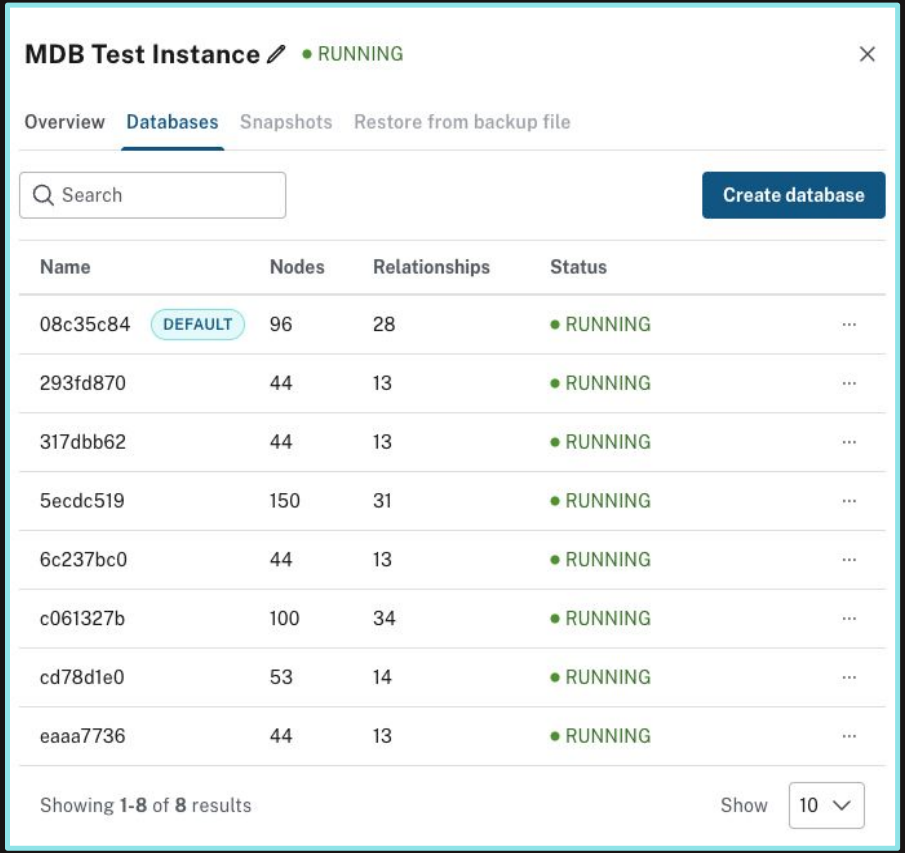

Virtual Graph DB

Neo4j Databases

Introducing Neo4j Virtual Graph Engine

A virtual graph database that brings full graph intelligence to your databases and lakehouses

Query your Databricks, Snowflake, or SQL DBs as a graph → with the complete Neo4j stack.

True Zero-Copy

No hidden caches. No background data duplication.

Your warehouse is the only system of record — always. Unlike competitors, we're explicit about where data lives.

60+ Graph Algorithms

PageRank, Louvain community detection, node similarity, and more via GDS Sessions.

Run real graph analytics — not just traversals.

AI-Powered Queries

Don't know Cypher? Copilot translates natural language to graph queries.

No learning curve required — ask questions in plain English.

Deploy to Value: Days

First query same day. Production insights in weeks, not quarters.

Start proving ROI before any major investment.

BEYOND VIRTUALIZATION → FULL GRAPH INTELLIGENCE

Anyone can query graphs. Graph Engine delivers the complete intelligence stack — algorithms, visualization, AI copilot — that turns a query into business impact.

Demo
Virtual Graph

Neo4j Databases

NEW CAPABILITIES

Virtual Graph

Aura(GA): Q4 2026

Self Managed (GA): Q1 2027

Zero copy graph querying on enterprise data sources.

- Supports: Databricks, BigQuery, Snowflake
- Schema -> Mapping -> VG
- Usable in all surfaces
- Limited Cypher

- Q3: improved Cypher surface, QPP
- Q3: Self managed experience
- Q4: Iceberg, AWS
- Q4: Caching/ingestion

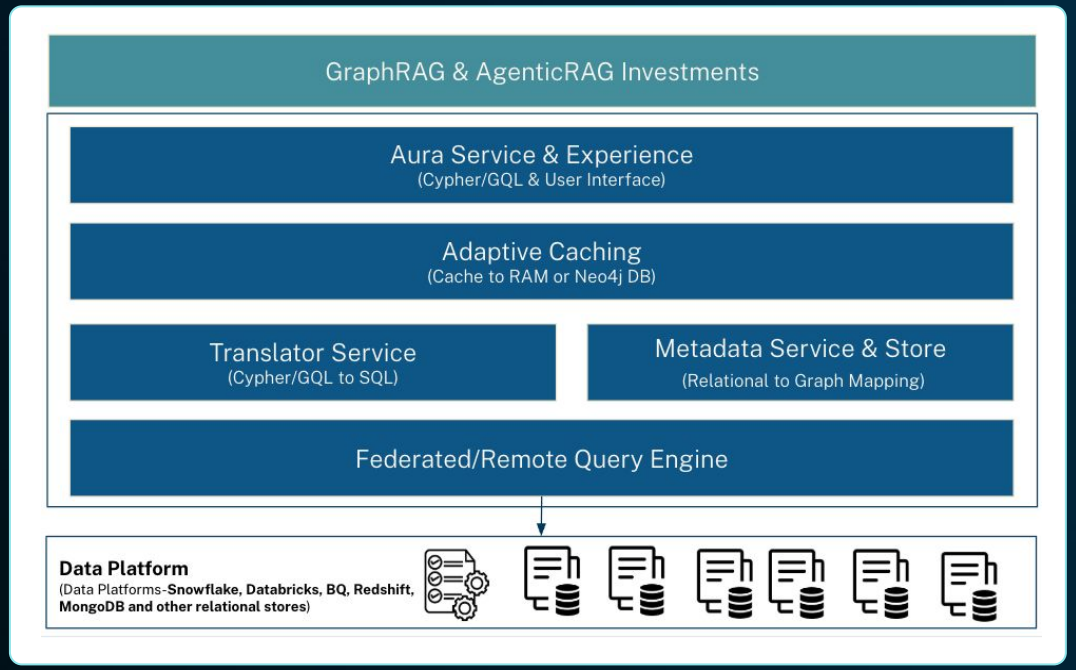

Neo4j Databases

Comparison: Native vs. Virtual Graph

Native Neo4j

⚡ REAL-TIME OLTP

Workloads: Transactional applications, high-frequency operational updates, real-time apps with sub-second SLAs.

Performance: Ultra-low latency (sub-100 millisecond) & ultra-high query throughput.

Agentic AI: Real-time AI agents requiring millisecond context retrieval and graph traversals.

Storage: Dedicated, index-free adjacency graph engine built from the ground up.

Best For: Context engine for agents, Fraud detection, real-time recommendations, identity resolution.

Virtual Graph

📈 ANALYTICS & FAST POC

Workloads: High-scale analytical workloads (OLAP), batch graph algorithms, exploratory query.

Performance: Scalable cloud warehouse pushdown compute for long-running AI agents.

Lakehouse GraphRAG: Grounding LLMs over existing Databricks, Snowflake & BigQuery data.

Rapid Proof-of-Value: Run graph PoCs in hours without building weeks of ETL pipelines.

PoC Accelerator: Replaces weeks of heavy ETL with an instant zero-copy trial, long running analytical agents

Neo4j Databases

Neo4j Deployment & Memory Sizing Guide

Recommended RAM-to-Data Ratios tailored by workload profile and deployment model

Real-Time Transactional

e.g., Fraud Detection, Identity, Payments

1:2 – 1:6
RAM TO DISK RATIO

- **Ultra-Low Latency:** Sub-millisecond lookup requirements mandate warm graph caching.
- **High Concurrency:** Active transactional writes and concurrent OLTP traversals.

High memory density for graph cache

Graph Analytical

e.g., Supply Chain, Impact & Lineage

1:12
RAM TO DISK RATIO

- **Deep Traversals:** Optimized for Graph Data Science (GDS) algorithms and pathfinding.
- **Batch Processing:** Balanced memory footprint for global multi-hop queries.

Optimized for operational analytics

GenAI & Vector

e.g., GraphRAG, Semantic Search

1:16
RAM TO DISK RATIO

- **HNSW Indexing:** Holds vector embeddings and graph grounding context efficiently.
- **AI Grounding:** Hybrid retrieval combining vector similarity and graph topology.

Efficient vector memory footprint

Virtual Graph

e.g., Databricks, Snowflake, BigQuery

1:32 – 1:50
RAM TO PLATFORM DATA RATIO

- **SQL Pushdown:** Offloads heavy compute directly to the customer's cloud data warehouse.
- **Lean Infrastructure:** Minimal memory footprint required for query translation & metadata.

Maximum cost efficiency per TB

Composite DB

Neo4j Databases

NEW CAPABILITIES

Composite Database

Self managed: Available

Aura: Q4 26 - H1 27

- **Unified Cypher Access**: Logical metadata layer that allows a single Cypher query to retrieve and join data across multiple constituent databases without storing graph data itself.
- **Data Federation & Sharding**: Supports both data federation and horizontal sharding.
- **Distributed Routing via Aliases**: Uses aliases to route sub-queries across separate instances or clusters.
- **Autonomous Domain Governance**: Preserves independent access controls, security policies, and resource management for each constituent database while exposing a single logical endpoint to applications.
- Available in self managed and coming soon to Aura.

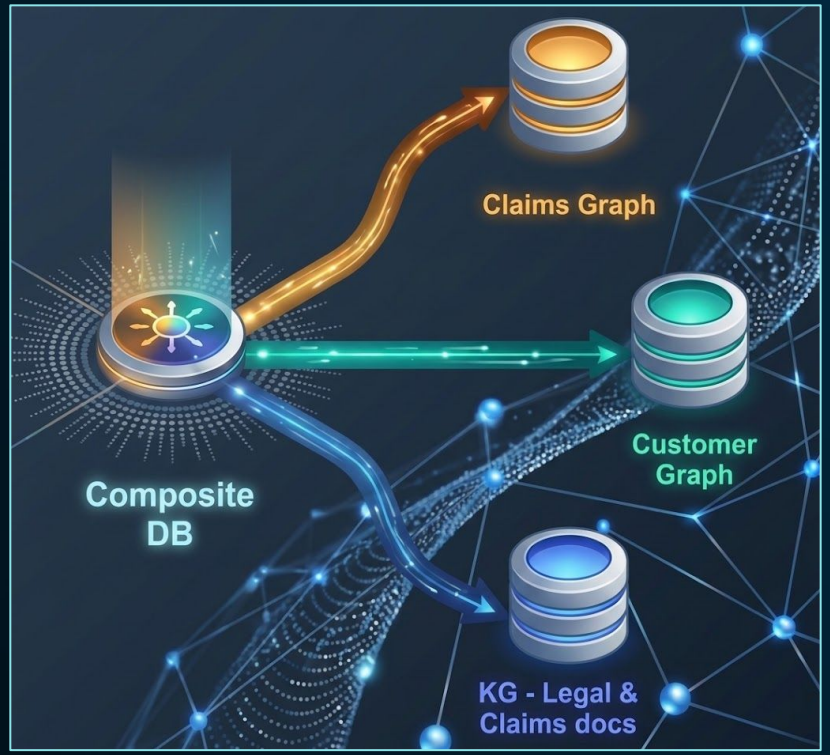

Security & Compliance

Neo4j Databases

CORE DATABASE ENHANCEMENTS

Security & Operations

Graph Type: Concise way to define an "open" schema. Unique relationship property, relationship key, property data types

Property Based Access Control (PBAC) for RBAC, data-driven rules to control READ and TRAVERSE privileges on nodes

AuraDB: Stronger login protection with MFA, **continuous monitoring** through real-time Security Logs forwarded to your cloud provider, activity feed tracking user actions

Controlled access using IP Filtering to restrict connections to trusted networks

Graph Type

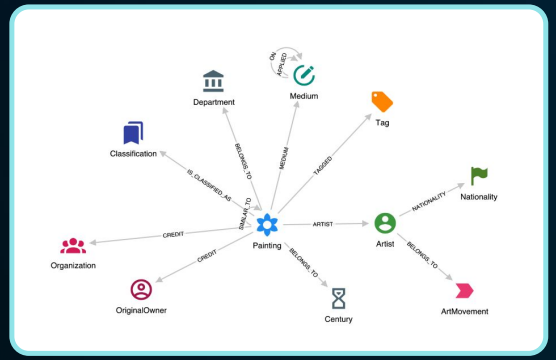

Property Based Access Control (PBAC)

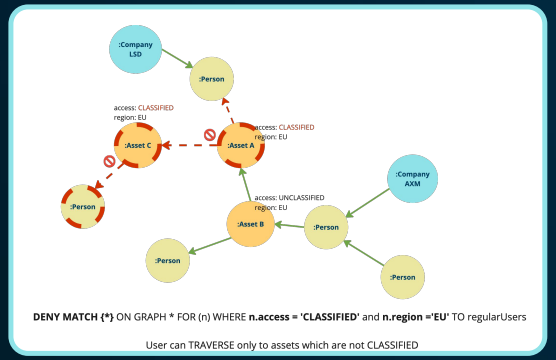

AuraDB: Security & Activity Logs & MFA

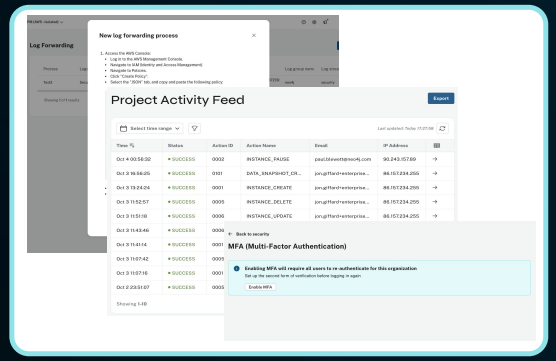

AuraDB: IP Filtering

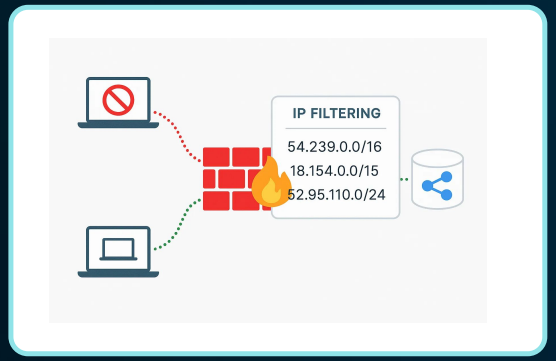

2024-2025

Neo4j Databases

CORE DATABASE ROADMAP

ABAC - Attribute-Based Access Control

Now Available

Dynamically control access based on attributes, such as IP addresses, time of day, or other custom environmental & contextual factors.

Compare user and data property in a PBAC rule (Oct ‘26)

Write custom plugins to apply metadata tags to users (Nov ‘26)

ABAC list intersection of user tags and data property (Dec ‘26)

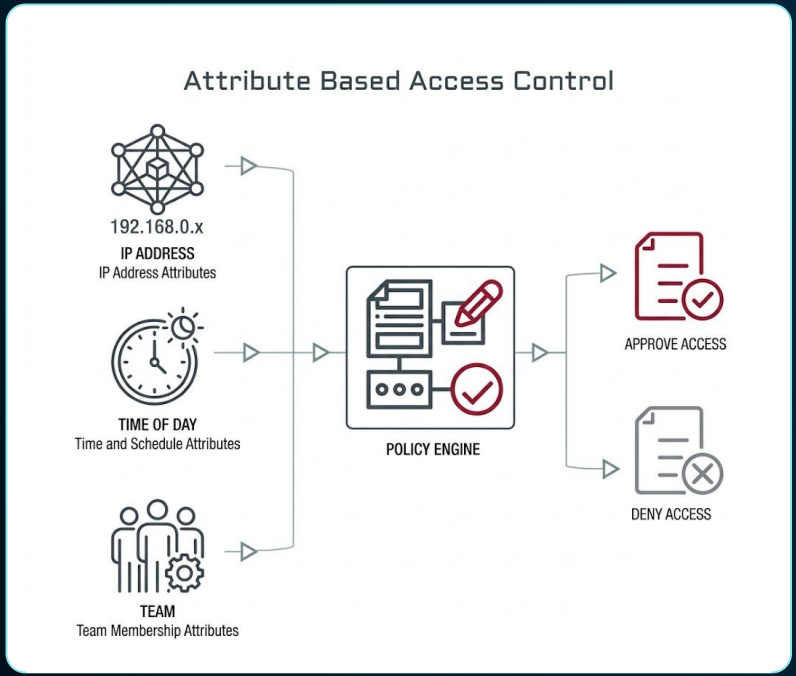

AI Powered Graph Tools

5 Seconds to Sign Up
Minutes to Wow
Days to Value

GraphAware Solution

Context Engine

- Ontology
- Studio
- Build Services
- Runtime
- Memory

Graph AI

- Analytics (Algorithms)
- Aura Agents
  Build multi hop agents in seconds.
- Agentic Ecosystem
  Integrate: AWS, GCP, Azure, LangChain, ServiceNow, SalesForce, Databricks & Snowflake

AI Powered Graph Tools

- Query
- Bloom
- Dashboards
- Data Import
- Doc Intel
- Ontology Mgmt

Neo4j Databases

- Enterprise
  Enterprise Edition
- InfiniGraph
- Composite DB

AI Powered Graph Tools

Studio: Unified Experience

Improved workflows with a single hub for all data management tasks

Consistent user experience through a **single unified console** across Neo4j tools

Easier collaboration as teams can share resources and collaborate on projects

Improved productivity with GenAI copilot by helping developers write and improve Cypher queries

Secure data access with expanded roles and new access controls

Data Import (growing sources)

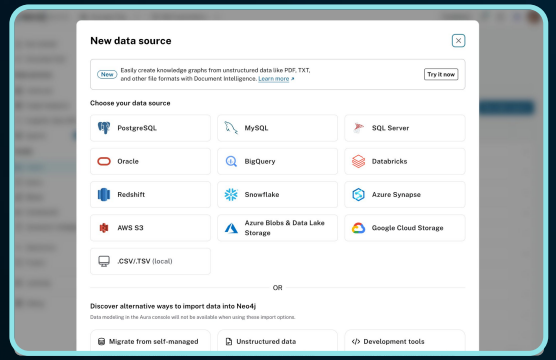

AI Powered Graph Data Modeling

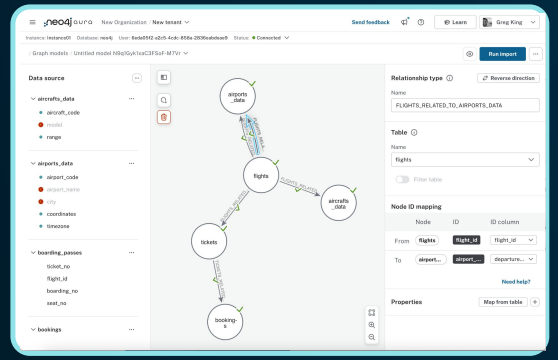

AI Powered Co-pilot for Cypher Adoption

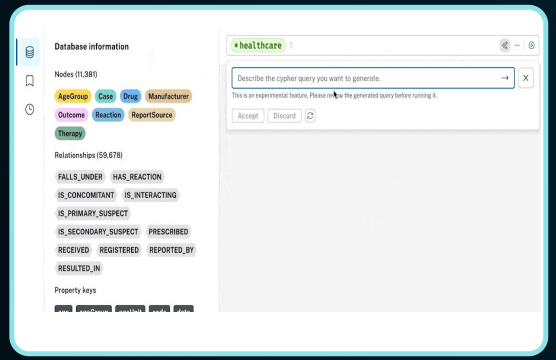

AI Powered Dashboards

Available

AI Powered Graph Tools

Neo4j Studio - Across all Form factors

Single Developer
(Free)

Neo4j Desktop Studio

query
bloom
import
dashboards

- **Local** or **remote** databases
- For **individual** testing and evaluation, **no sharing**
- **AI assistant**
- Import from **files only**

Basic functionality, free for everyone

Enterprise Collaboration - Cloud
(Free, included with Aura)

Neo4j Aura Studio

query
bloom
agents
import
dashboards
and more

- **Aura** databases
- **Self-managed** databases (via Fleet Manager)
- **Multi-project** collaboration
- **Cloud storage** for tools data
- **AI assistant** and **agents**
- **Cloud** and **files** imports

Complete & fully-managed experience

Enterprise Collaboration - Self-managed
(Paid product)

Neo4j Enterprise Studio

query
bloom
dashboards

Available

- **Self-managed** databases only
- **Single project** collaboration
- **Self-managed** storage for tool data

Graph tools for self-managed customers

AI Powered Graph Tools

NEW CAPABILITIES

Document Intelligence

GA: Oct 2026

VDC: Q4 2026

Use LLMs with a graph schema to construct high quality knowledge graphs efficiently from unstructured data

- Integrated experience in Aura console
- Document upload or cloud buckets
- Schema: AI generated from sampled docs
- Interactive, contextual schema assistant
- Processes 100s of pages
- No Self Managed experience (yet)

- Oct '26 - GA Self service: scalability, 1000s of pages, complex layouts, extraction quality
- Q4: Enterprise, VDC, Model Choice, Scale, API/CLI

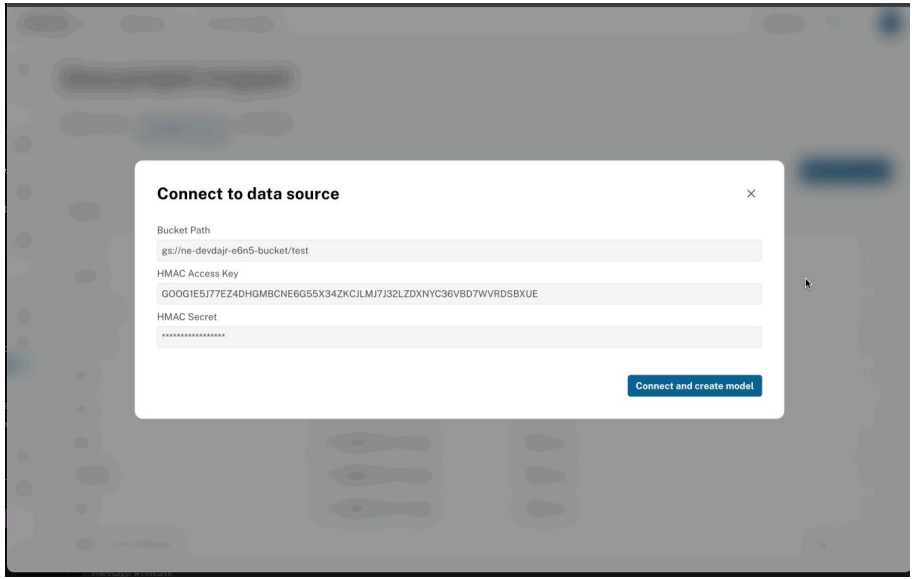

Product Announcement

Ontology for Agentic AI

Moving beyond passive database management. Neo4j is introducing an AI-native ontology layer designed as cognitive scaffolding for autonomous LLM agents to reason over interconnected enterprise data with provable accuracy.

Agentic Reasoning | Deterministic Truth | MCP-Native Integration

AI Powered Graph Tools

Agent-Optimized Capabilities

Semantic Graph Schema

- **Nodes, Labels & Relationships:** Complete typing for node labels, relationship types, and scalar/vector properties for deterministic query compilation.
- **LLM Semantic Descriptions:** Docstring metadata attached directly to labels, relationships, and attributes to guide LLM prompt context and disambiguation.
- **Cardinality & Property Enums:** Explicit relational bounds (1:1, 1:N, N:M) and strict enum values that eliminate hallucinated relationships and states.
- **Graph Integrity Constraints:** Built-in structural rules and schema validations ensuring agents never execute illegal graph mutations.

Agentic Reasoning & Tools

MCP NATIVE

- **Direct MCP Tool Mapping:** GraphRAG retrievers map automatically into Model Context Protocol tools for autonomous tool-use.
- **Multi-Hop Reasoning:** Equips LLMs to traverse interconnected paths across entities, solving multi-faceted queries where single-lookup RAG fails.
- **High-Precision Text2Cypher:** Enriched semantic context allows agents to synthesize accurate Cypher queries grounded in exact domain ontology.
- **Traceable Agent Memory:** Anchors long term and reasoning memory to verifiable subgraphs, eliminating context drift and guaranteeing auditability.

Demo
Ontology Capabilities

AI Powered Graph Tools

MOST DEVELOPER-FRIENDLY PLATFORM FOR AGENTIC AI

Neo4j Ontology

Ontology Spec: Publish the Neo4j Ontology Spec to support interoperability with formats like RDF and JSON-LD.

Ontology Persistence: Store the database ontology with the database with a Cypher surface for CRUD operations and lifecycle management.

Ontology Management: Develop comprehensive management tooling for ontologies across the API, UI, CLI and MCP surfaces, including inferring ontologies from data, importing from third-party sources and providing reference domain ontologies by use case.1

Ontology-aware Tools: Integrate first-class ontology awareness across the Neo4j stack to enhance all workflows like Onboarding assistance, Agent/MCP tool configurations, Data ingestion, Virtual graph setup, Text2Cypher generation and more.

Coming soon

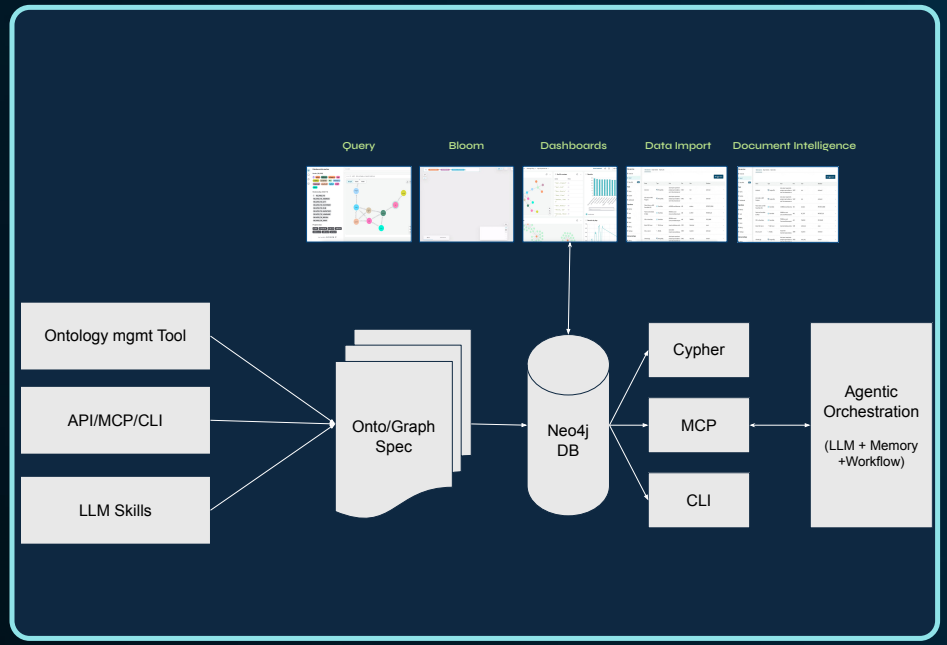

Graph AI

Enterprise-grade graph algorithms & agents that uncover hidden patterns and predictive insights.

GraphAware Solution

- Intelligence Analytics
- Financial Crime

Context Engine

- Ontology
- Studio
- Build Services
- Runtime
- Memory

Graph AI

- Analytics (Algorithms)
- Aura Agents
  Build multi hop agents in seconds
- Agentic Ecosystem
  Integrate: AWS, GCP, Azure, LangChain, ServiceNow, SalesForce, Databricks & Snowflake

AI Powered Graph Tools

- Query
- Bloom
- Dashboards
- Data Import
- Doc Intel
- Ontology Mgmt

Neo4j Databases

- Enterprise Edition
- InfiniGraph
- AuraDB
- Virtual DB
- Composite DB

Graph AI

GRAPH ANALYTICS AS KEY DIFFERENTIATOR

Graph Analytics

Graph algorithms now run on dedicated, on-demand machines with their own memory and CPU, so analytics never compete with your transactional workload.

Available Now — **Aura Graph Analytics**: pay-as-you-go and serverless, on any cloud

EAP September — **Neo4j Graph Analytics:** brings this architecture to self-managed Neo4j, on any cloud or on-prem

New — **Graph Embeddings API**: embeddings out of the box, with no model to train

Coming Soon — **No-code analytics jobs** with **AI-powered algorithm recommendations**

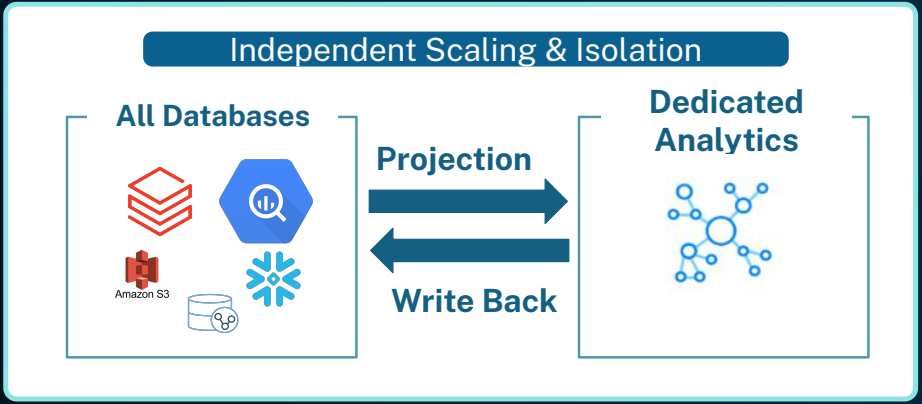

Improve decisions from connected data, beyond the limits of traditional analytics

Reduce time-to-value with 65+ built-in algorithms on any enterprise data

Lower TCO by sizing analytics compute to the job, not the database

Graph AI

MOST DEVELOPER-FRIENDLY PLATFORM FOR AGENTIC AI

Agentic Engineering with Neo4j

Discoverability
can agents find Neo4j context and capabilities

neo4j.com/llms.txt

- drivers
- MCP
- graphrag
- skills
- Modeling, cypher
- docs, graphacademy

Usability
can agents effectively run Neo4j skills and tools

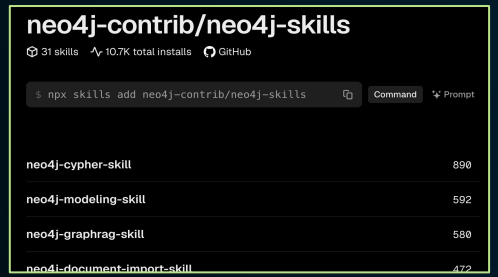

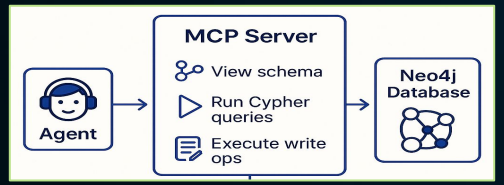

Invocability
can agents invoke Neo4j via CLI, API, and MCP

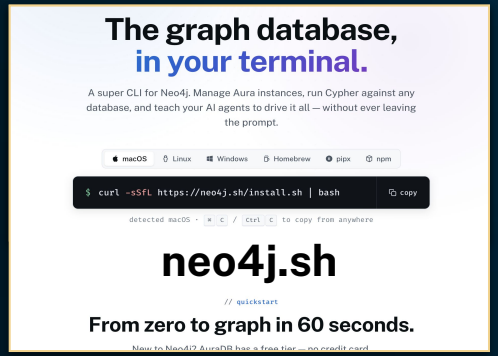

- consolidated API
- operations latency

AX measures how easily AI agents discover, use, and invoke Neo4j. As agents become primary API consumers, it requires deliberate investment across three pillars.

Graph AI

MOST DEVELOPER-FRIENDLY PLATFORM FOR AGENTIC AI

MCP Servers

MCP is the AI native tool integration protocol for agents.

For APIs, Databases, SaaS services.

Official MCP Server for

- Core Database & GDS (Q4 25)
- Aura Agents (Q1 26)
- **Hosted: Aura Databases (Q2 26)**
- Graph Analytics/AGA (Q3 26)
- Aura APIs (Aura mgmt) (H2 26)

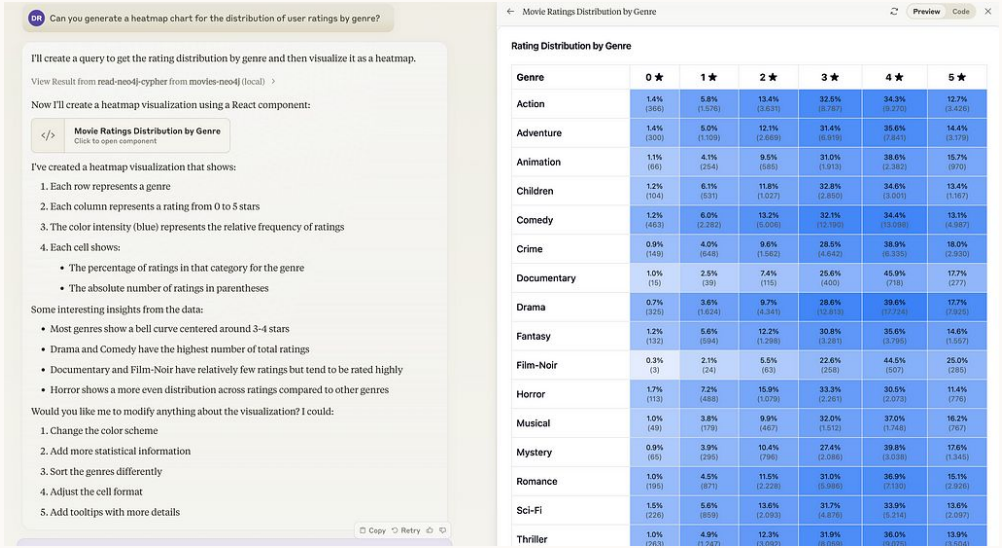

Genre	0 ★	1 ★	2 ★	3 ★	4 ★	5 ★
Action	1.4% (366)	5.8% (1.576)	13.4% (3.631)	32.6% (8.787)	34.3% (9.270)	12.7% (3.426)
Adventure	1.4% (300)	5.0% (1.109)	12.1% (2.669)	31.4% (6.919)	35.6% (7.841)	14.4% (3.179)
Animation	1.1% (66)	4.1% (254)	9.5% (585)	31.0% (1.913)	38.6% (2.382)	15.7% (970)
Children	1.2% (104)	6.1% (531)	11.8% (1.027)	32.8% (2.850)	34.6% (3.001)	13.4% (1.167)
Comedy	1.2% (463)	6.0% (2.282)	13.2% (5.008)	32.1% (12.190)	34.4% (13.098)	13.1% (4.987)
Crime	0.9% (149)	4.0% (648)	9.6% (1.562)	28.5% (4.642)	38.9% (6.335)	18.0% (2.930)
Documentary	1.0% (15)	2.5% (39)	7.4% (115)	25.6% (400)	45.9% (718)	17.7% (277)
Drama	0.7% (325)	3.6% (1.624)	9.7% (4.341)	28.6% (12.813)	39.6% (17.724)	17.7% (7.925)
Fantasy	1.2% (132)	5.6% (594)	12.2% (1.298)	30.8% (3.281)	35.6% (3.795)	14.6% (1.557)
Film-Noir	0.3% (3)	2.1% (24)	5.5% (63)	22.6% (258)	44.5% (507)	25.0% (285)
Horror	1.7% (113)	7.2% (488)	15.9% (1.079)	33.3% (2.261)	30.5% (2.073)	11.4% (776)
Musical	1.0% (49)	3.8% (179)	9.9% (467)	32.0% (1.512)	37.0% (1.748)	16.2% (767)
Mystery	0.9% (65)	3.9% (295)	10.4% (796)	27.4% (2.086)	39.8% (3.038)	17.6% (1.345)
Romance	1.0% (195)	4.5% (871)	11.5% (2.228)	31.0% (5.986)	36.9% (7.130)	15.1% (2.926)
Sci-Fi	1.5% (226)	5.6% (859)	13.6% (2.093)	31.7% (4.876)	33.9% (5.214)	13.6% (2.097)
Thriller	1.0% (263)	4.9% (1.247)	12.3% (3.082)	31.9% (8.089)	36.0% (9.075)	13.9% (3.504)

Now Available

Graph AI

MOST DEVELOPER-FRIENDLY PLATFORM FOR AGENTIC AI

Aura Agent

- Available Capabilities
  - Low code and AI assisted agent **management**
  - GraphRAG **tools**: Cypher, Vector + Cypher Search, Text2Cypher
  - Agent Test: Iterate on your agent with built-in chat
  - **API**: Aura API REST endpoint, MCP + M2M Auth
  - Agent Evaluation: Score your agent performance
- Pricing $0.35 per agent hour for API or MCP Server endpoint access
- Available in **tiers**: Aura Free, Professional and Business Critical
- Roadmap
  - Conversation History (Q4 2026)
  - Agents for **Enterprise** (AWS/GCP) (Q1 2027)

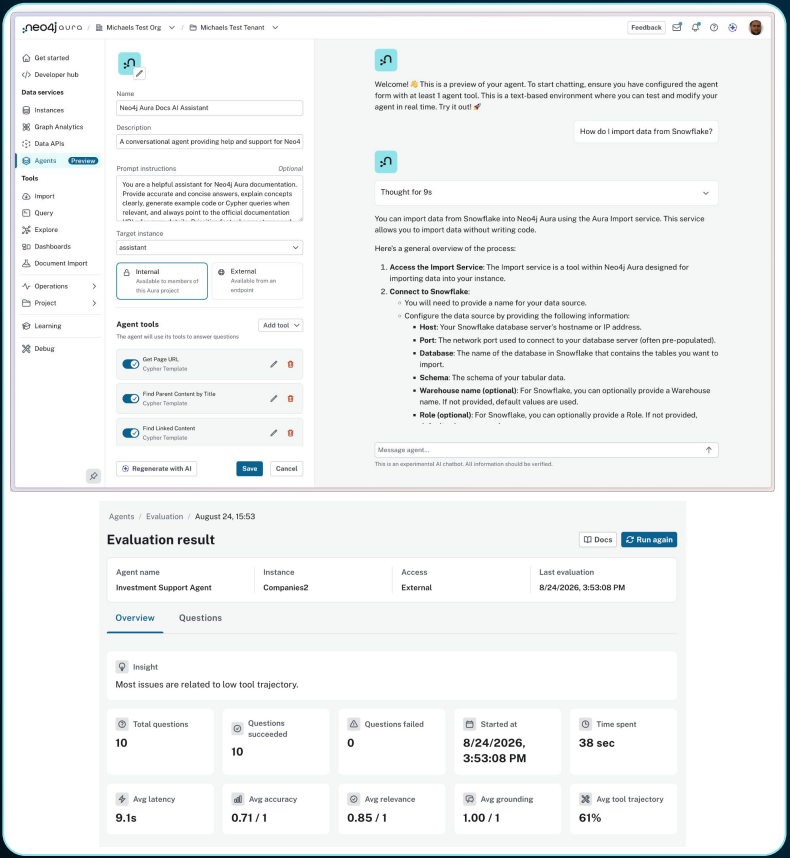

Evals - Available Now

Demo

Aura Agent - Evals

Graph AI

Deep integrations across the AI ecosystem

AI & Agent Ecosystem

Integrated with AI ecosystem through MCP incl. Agent Memory

- AI Libraries & Frameworks
- Agent Frameworks
- CSP AI platforms
- Enterprise Agent Platforms

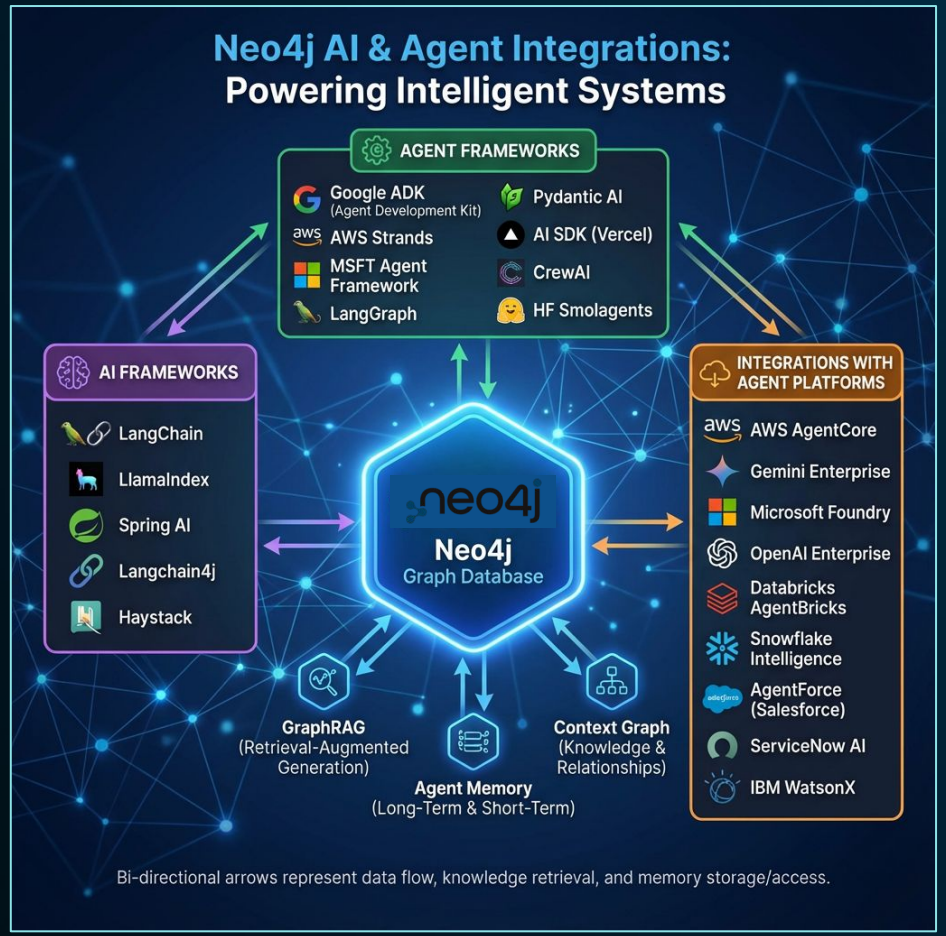

Ongoing

Neo4j Database Offerings

Open, Scalable, Secure

GraphAware Solution

- Intelligence Analytics
- Financial Crime

Context Engine

- Ontology
- Studio
- Build Services
- Runtime
- Memory

Graph AI

- Analytics (Algorithms)
- Aura Agents
  Build multi hop agents in seconds.
- Agentic Ecosystem
  Integrate: AWS, GCP, Azure, LangChain, ServiceNow, SalesForce, Databricks & Snowflake

AI Powered Graph Tools

- Query
- Bloom
- Dashboards
- Data Import
- Doc Intel
- Ontology Mgmt

Neo4j Databases

- Enterprise Edition
- InfiniGraph
- AuraDB
- Virtual DB
- Composite DB

Context Engine

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.

Learning

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

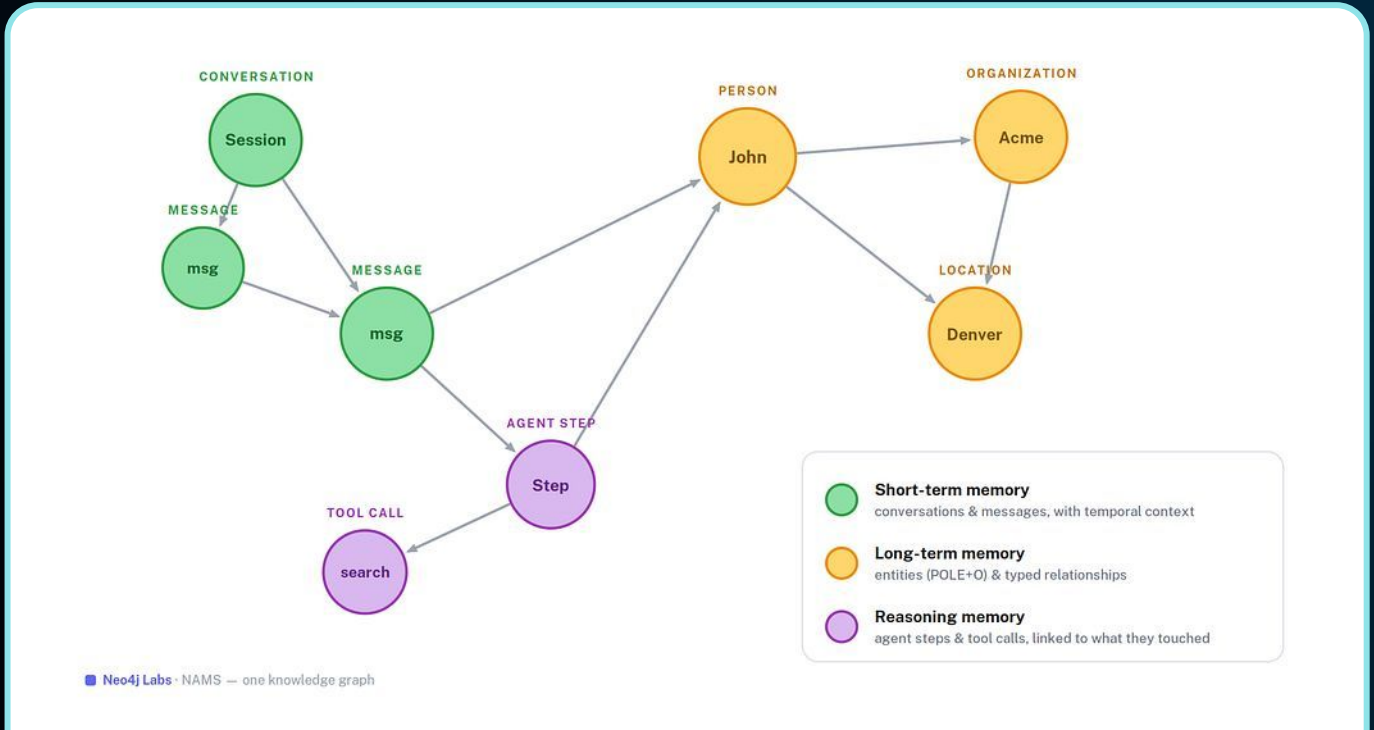

Learning

From conversation to knowledge to action

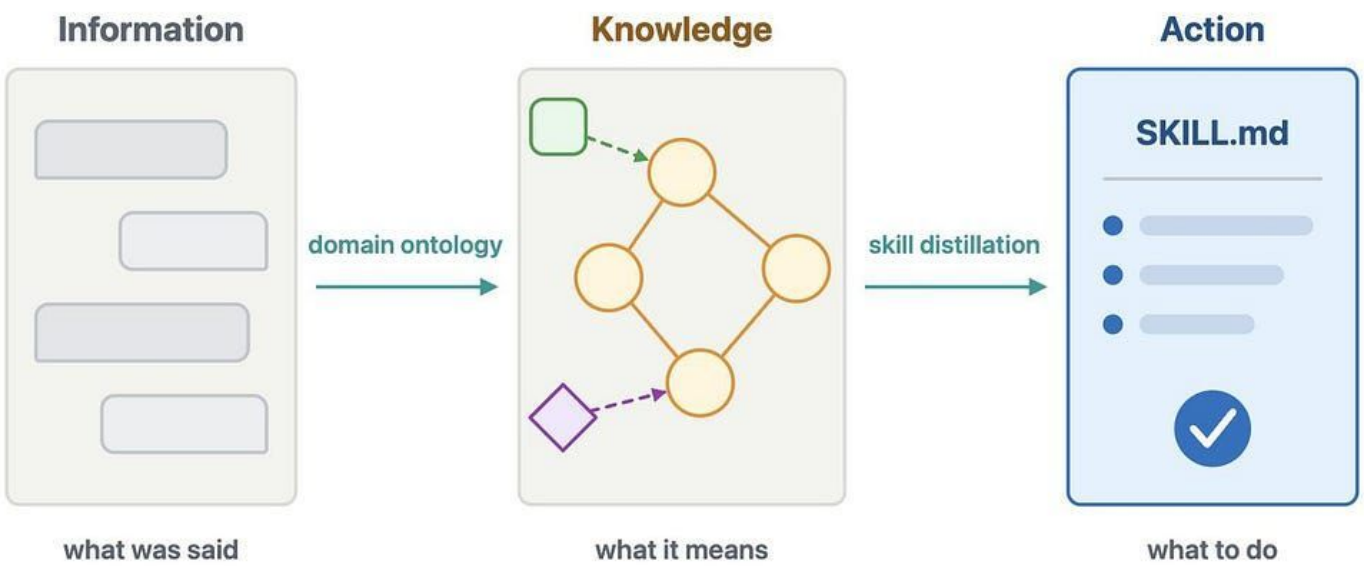

Demo
Memory Service

Neo4j's Graph Intelligence Platform

The most comprehensive platform to build an enterprise knowledge layer

GraphAware Solution

- Intelligence Analysis
- Financial Crime

Context Engine

- Ontology
- Studio
- Build Services
- Runtime
- Memory

Graph AI

- Analytics (Algorithms)
- Aura Agents
  Build multi hop agents in seconds.
- Agentic Ecosystem
  Integrate: AWS, GCP, Azure, LangChain, ServiceNow, SalesForce, Databricks & Snowflake

AI Powered Graph Tools

- Query
- Bloom
- Dashboards
- Data Import
- Doc Intel
- Ontology Mgmt

Neo4j Databases

- Enterprise Edition
- InfiniGraph
- AuraDB
- Virtual DB
- Composite DB

Thank you to our customers

airbnb
AIRBUS
Albertsons
AstraZeneca
die Bayerische
BNP PARIBAS PERSONAL FINANCE
Boston Scientific
CATERPILLAR
CenturyLink
CISCO
citibank
COMCAST
CRÉDIT AGRICOLE
DAIMLER
DB BAHN
Deloitte.
Discovery
El Corte Inglés
FUJITSU
intuit
LEVI STRAUSS & CO.
LOCKHEED MARTIN
lyft
MITSUBISHI
Microsoft
NASA
NORDSTROM
NOVARTIS
novo nordisk
pwc
pitney bowes
SAMSUNG
Santander
Staples
Thomson Reuters
UBS
U.S. ARMY
verizon
Walmart
Western Union
ZURICH