



Shaping the future of manufacturing with connected data

GraphTalk Manufacturing/Automotive 2026



Top AI/Graph Use Case Categories in Manufacturing and Automotive

Top performing projects in 2025/2026



Data & Shop Floor visibility

The Goal:

- **Global equipment visibility:** a single overview of all installed machines and lines across the network, which directly supports relocation and investment decisions (i.e., "do we actually need to buy a new machine, or is one idle in another plant?")
- **Standardized, reusable access:** data is exposed via API in a consistent format, so multiple tools and use cases can consume it rather than each building its own integration
- **Centralized data quality tracking:** even though data originates locally at each plant, quality can be monitored centrally across the whole distributed network



Supply Chain and Order Optimization



The Goal:

- **Multi-hop supplier reasoning:** "if Supplier A is delayed, which alternate suppliers also carry this part, and which of *those* are already qualified for this plant?"
- **Bill-of-materials cascade effects:** a graph connecting part → sub-assembly → finished product → demand forecast lets the agent see that reordering Part X should also flag Parts Y and Z
- **Cross-plant idle-stock visibility:** before ordering new stock, the agent should traverse the graph to check if another plant in the network already has surplus of that part, avoiding a purchase entirely rather than just optimizing the purchase.

Requirements Analysis and Product Engineering

The Goal:

- **Traceability & impact analysis:** Model requirement → design element → test case → verification result as connected nodes.
- **Compliance mapping:** Link requirements to the specific regulatory clause (e.g., a DO-178C or ISO 26262 paragraph) they satisfy. An auditor's question "show me evidence for clause 4.2" becomes a graph query with a citable path, not a document search.
- **Duplicate & conflict detection:** Represent similar requirements across programs/platforms as nodes; graph structure plus embeddings can surface near-duplicate or contradictory requirements





Sales & Marketing Optimization




The Goal:

- **Automated tender compliance drafting:** Model tender requirement → past project reference → certification → product capability as connected nodes. manually digging through old submissions.
- **Faster, more consistent win-theme reuse:** Once a proof point (a past delivery, a certification, a benchmark result) is in the graph, every future tender referencing a similar requirement can reuse it automatically.
- **Explainable recommendations:** Represent customer → purchased product → product relationship (complementary, bundle, successor) → industry segment as graph edges. A recommendation like "customers in this segment who bought A also adopted B within 6 months" i



The Challenge

The bottleneck is context engineering,
not model intelligence



The industry evidence (Siemens Research*)



31–94%

Of Data is siloed
and partly
governed but
differ in
semantics and
definitions

10→66%

*Hallucination
jump when a
model receives
insufficient
context

>40%

Of AI projects
may be
abandoned
without AI-ready
data



“When the context is not clear, AI projects will fail.”
-Voith





How context breaks down (Top 3 Reasons)

Siloed disconnection

Every type of data belongs to their prior defined toolset.
Snowflake, Databricks, Salesforce, Confluence, SAPs etc.

Context fragmentation


Different tools and sources feed competing definitions of the same entity, creating duplication and weak alignment.

Unstructured treasure

Roughly 85% of enterprise data is unstructured. Most valuable context is not sitting in clean tables

AI ready by Graph Design

Semantics and Relationships matters



The Problem aka The usual look:




	A	B	C	D	E	F	G	H
1	[IDENTITY] Node_ID	[IDENTITY] Node_Type	[IDENTITY] Name	[IDENTITY] Status	[CLASSIFICATION] Category	[CLASSIFICATION] Criticality	[CLASSIFICATION] Asset_Type	[CLASSIFICATION] Compliance_Fraser
2	SVC001	Service	Email & Collaboration	Active	Communication	High		
3	SVC002	Service	ERP Finance	Active	Business	Critical		
4	SVC003	Service	Network Connectivity	Active	Infrastructure	Critical		
5	SVC004	Service	Identity & Access Management	Active	Security	Critical		
6	SVC005	Service	Data Warehouse	Active	Analytics	High		
7	SVC006	Service	Customer Portal	Active	Business	High		
8	SVC007	Service	HR Self-Service	Active	Business	Medium		
9	SVC008	Service	Monitoring & Observability	Active	Infrastructure	High		
10	APP001	Application	Microsoft 365	Active		High		
11	APP002	Application	SAP ERP	Active		Critical		
12	APP003	Application	Cisco ACI	Active		Critical		
13	APP004	Application	Azure Active Directory	Active		Critical		
14	APP005	Application	Snowflake	Active		High		
15	APP006	Application	Salesforce	Active		High		
16	APP007	Application	Workday	Active		Medium		
17	APP008	Application	Datadog	Active		High		
18	APP009	Application	SAP BW	End of Life		Medium		
19	APP010	Application	Active Directory On-Prem	Active		Critical		
20	AST001	Asset	Dell PowerEdge R750 #1	Active			Hardware	
21	AST002	Asset	Dell PowerEdge R750 #2	Active			Hardware	
22	AST003	Asset	Cisco Nexus 9300 Core-01	Active			Hardware	
23	AST004	Asset	Cisco Nexus 9300 Core-02	Active			Hardware	
24	AST005	Asset	HPE ProLiant DL380 #1	Active			Hardware	
25	AST006	Asset	Microsoft 365 E3 License	Active			Software	
26	AST007	Asset	SAP S/4HANA License	Active			Software	
27	AST008	Asset	Snowflake Enterprise License	Active			Software	
28	AST009	Asset	Datadog Pro License	Active			Software	
29	AST010	Asset	Azure AD P2 License	Active			Software	
30	AST011	Asset	svc-erp-admin@corp.com	Active			Account	
31	AST012	Asset	svc-monitoring@corp.com	Active			Account	
32	AST013	Asset	admin@corp.onmicrosoft.com	Active			Account	
33	AST014	Asset	Cisco Firepower 2140	Active			Hardware	
34	AST015	Asset	Workday Enterprise License	Active			Software	
35	PRV001	Provider	Microsoft	Active	Software & Cloud			
36	PRV002	Provider	SAP	Active	Software			
37	PRV003	Provider	Cisco	Active	Hardware & Software			

The Problem aka The usual look:



[IDENTITY] Node_ID	[IDENTITY] Node_Type	[IDENTITY] Name	[IDENTITY] Status	[CLASSIFICATION] Category	[CLASSIFICATION] Criticality	[CLASSIFICATION] Type	[CLASSIFICATION] Compliance
SVC001	Service	Email & Collaboration	Active	Communication	High		
SVC002	Service	ERP Finance	Active	Business	Critical		
SVC003	Service	Network Connectivity	Active	Infrastructure	Critical		
SVC004	Service	Identity & Access Management	Active	Security	Critical		
SVC005	Service	Data Warehouse	Active	Analytics	High		
SVC006	Service	Customer Portal	Active	Business	High		
SVC007	Service	HR Self-Service	Active	Business	Medium		
SVC008	Service	Monitoring & Observability	Active	Infrastructure	High		
APP001	Application	Microsoft 365	Active		High		
APP002	Application	SAP ERP	Active		Critical		
APP003	Application	Cisco ACI	Active		Critical		
APP004	Application	Azure Active Directory	Active		Critical		
APP005	Application	Snowflake	Active		High		
APP006	Application	Salesforce	Active		High		
APP007	Application	Workday	Active		Medium		
APP008	Application	Log	Active		High		
APP009	Application	End of Life			Medium		
APP010	Application	On-Prem	Active		Critical		
AST001	Asset		Active			Hardware	
AST002	Asset		Active			Hardware	
AST003	Asset		Active			Hardware	
AST004	Asset		Active			Hardware	
AST005	Asset		Active			Hardware	
AST006	Asset	Microsoft 365 E3 License	Active			Software	
AST007	Asset	License	Active			Software	
AST008	Asset	License	Active			Software	
AST009	Asset		Active				
AST010	Asset		Active				
AST011	Asset	svc-mon	Active				
AST012	Asset	admin@corp.com	Active				
AST013	Asset	Cisco Firepower	Active				
AST014	Asset	Workday Enterprise License	Active				
AST015	Asset	Microsoft	Active	Software & Cloud			
PRV002	Provider	SAP	Active	Software			
PRV003	Provider	Cisco	Active	Hardware & Software			

How it should look like:






Graph Analytics

Communities
Clusters
Entity Resolution

Graph

Superset data structure for retrieval



[< >]

Agents


Query Templates
Hybrid Search
Query Generation
Memory

Standard ETL



Vectors

Embedder



Extracted
Entities

LLM



Structured

(Relational, flat files, etc.)



Semi-Structured

(NoSQL, document, etc.)

Unstructured
(text, images, etc.)



AI & KG Architecture



GenAI Applications

LLM Orchestration inc. retrievers, tools & agents

Knowledge Graph



mapping unstructured + structured data

LLM Platform



Models & Reasoning Engine

Enterprise Data Platform & External Databases




Knowledge Graphs Unlock Quality AI

KGs

Knowledge

Facts

Context



GenAI

Language

Statistics

Creativity

Knowledge graphs = AI-ready data



Contextual

Relationships are treated as primary data structures

Entities and their connections are built into the data model

Flexible

Schema grows & adapts to meet evolving business needs

Evolve your model while queries keep running

Standardized

Expressive data model provides a common language

Consistent meaning across all your data sources

Connected data. Smarter AI.

230% ROI.

As reported in IDC White Paper, The Business Value of Neo4j Graph Intelligence Platform*

Get the whitepaper →

\$4.16M

average annual benefit

230%


three-year return on investment

7.8 months

payback on investment

<https://neo4j.com/the-business-value-of-neo4j-graph-intelligence-platform/>

One Particular Example - Voith Virtual Graphs



Voith - Virtual Graph Layer

Base layer

- Single source of truth for all entities and their relationships of a single tool in a graph data product
- Built once, maintained centrally — no duplication across teams
- Human-readable and verifiable, so domain experts can validate and correct it

Virtual layers

- Each department, customer, or use case gets its own constrained view or copy of the base graph data product
- Complexity is reduced per audience: operators see what operators need, engineers see what engineers need
- Isolate or merge base graph products as needed — context stays clean across silos






Neo4j's Virtual Graph is the answer, but to **what** question ?

“How can my data and AI team leverage the expressive power of graph databases, while maintaining a single, unmoved source of truth within my existing data storage infrastructure?”




How it works





Enterprise Architecture for Virtual Graph




Next Steps

- Register for Virtual Graph Access —————→
- Read [Documentation](#) and [Blog](#)
- Deploy Neo4j on [Snowflake Marketplace](#)
- Read more about our partnership: neo4j.com/snowflake
- Send questions to ecosystem@neo4j.com
- Register for upcoming [partner ecosystem events](#)



Summary



Conclusion

**Organizations are
drowning in data but
AI is still starving for
context**



Start connecting siloed data with semantic relationships



Enhance retrieval AI-Systems with GraphRAG and utilize rich context



Keep a flexible Data architecture to keep pace with changing AI requirements