

Building the Context Graph

→ A Practitioner's Perspective

Foundation Capital recently described context graphs as AI's trillion-dollar opportunity. We have been building exactly this for years. Here is what we learned about turning enterprise knowledge into a living, queryable graph.

Content

Intro

01 The Thesis We Share

02 Where Our Paths Converge

03 Where We Go Further

04 What We Learned the Hard Way

05 An Invitation

→ References

Building the Context Graph: A Practitioner's Perspective

Author: Viktor Voss, Senior Software Architect

In December 2025, Jaya Gupta and Ashu Garg of Foundation Capital published „[AI's Trillion-Dollar Opportunity: Context Graphs](#)“. Their thesis: the next generation of enterprise platforms will not be built by adding AI to existing data, but by capturing the decision traces that make data actionable — the reasoning, exceptions, precedents, and cross-system context that explain why things happened, not just what happened.


I read it with interest. It describes, in the language of venture capital, what we have been building with Bardioc for years — not as a theoretical exercise, but as a production system processing hundreds of millions of API calls per day across hundreds of terabytes of semantically connected data.

This is our perspective as practitioners: what we agree with, where our paths converge, where we go further, and what we have learned that might add to the conversation.

01 The Thesis We Share

Foundation Capital's central insight: traditional systems of record — CRMs, ERPs, ticketing systems — store what happened but lose why it happened. The reasoning that connects data to action was never treated as data in the first place. They call the structure that captures this reasoning a context graph.

When we started building Bardioc, our initial observation was almost identical. Machines cannot compute causality on their own, but causality is essential for understanding the world. The way to give machines this ability is to represent relationships as machine-readable data — not isolated records but semantically connected knowledge: facts, their relationships, their temporal evolution, and the reasoning that ties them together.



02 Where Our Paths Converge

The parallels between Foundation Capital's context graph and Bardioc's architecture run deep. Three of them are worth naming explicitly.



Semantic integration across system boundaries

No incumbent system captures the cross-system synthesis that humans do every day — checking the CRM, reading a Slack thread, reviewing an escalation, making a judgment call. Bardioc was designed from the ground up as a semantic data integration platform. Our bidirectional connectors maintain live, synchronized links to source systems, not after-the-fact ETL extractions. When data changes in the source, the graph updates. When insights emerge in the graph, they can flow back. This is what Foundation Capital means by being “in the write path, not the read path”.



A universal ontology instead of ad-hoc schemas

Bardioc includes OGIT (Open Graph of IT) — an [open-source, modular ontology](#) that represents the world from the perspective of enterprises and their processes: people, organizations, business relationships, locations, IT systems, and the connections between them. OGIT evolves through a governed, community-driven process. Developers have maximum freedom to define their own data, with the incentive to formalize and share their schemas for broader reuse. A context graph that cannot evolve with the world it represents will calcify.

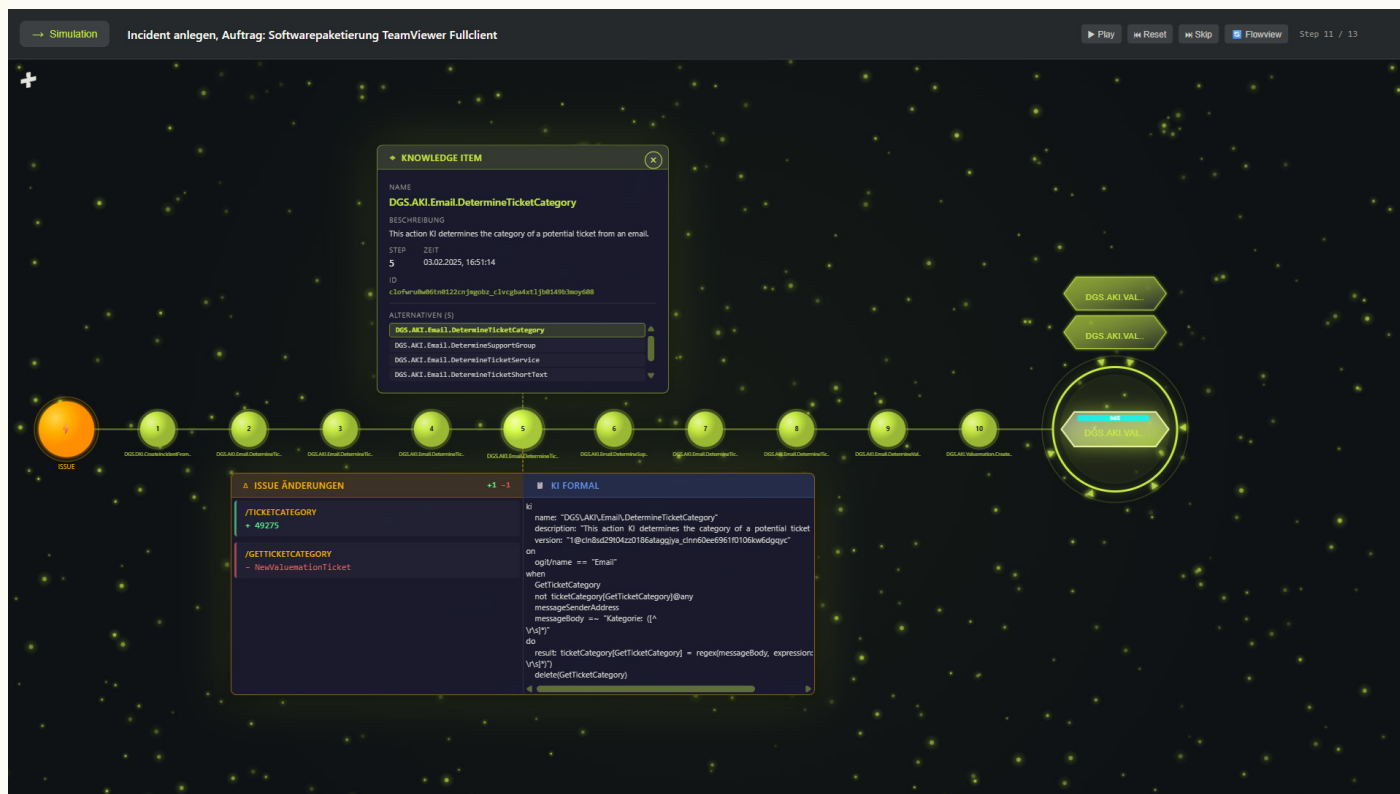


Capturing the “why”, not just the “what”

Foundation Capital's decision traces record how context turned into action: what inputs were gathered, what rules applied, what exceptions were granted, who approved. In Bardioc, the Reasoning Engine does not only record decision traces — it produces them. The Reasoning Engine is a hybrid, knowledge-based AI that works directly on the semantic graph. Every decision it makes is traceable: which knowledge building blocks contributed, which data was considered, what logical path led to the conclusion. With our Decision Path visualization, any stakeholder — not just engineers — can inspect exactly how and why a specific decision was reached.

These three layers — ontology, context graph, decision traces — are not three competing concepts. They are layers of the same architecture.

The 2026 industry consensus has settled there. [As Atlan put it](#) in their analysis: “the ontology provides the semantic foundation; the context graph adds the operational layer.” Another analysis was even sharper: “a context graph without ontological grounding works for prototypes but breaks in production.” This is why Bardioc looks the way it looks. It was built with the same conclusion in mind from the start.



The Decision Path visualization makes every step of the Reasoning Engine's logic transparent — from data inputs through knowledge building blocks to conclusion.

03 Where We Go Further

Our experience operating Bardioc has led us to extend the article's thesis in three directions.

One is the move from runtime traces to structured knowledge. The article describes context graphs as emerging from agent execution traces — agents orchestrate workflows, emit decision traces, and over time the traces accumulate into a queryable graph. That is one valid way to build the layer. It is not the only one. In Bardioc, the context graph is not a byproduct of agent execution. It is the primary asset, built intentionally through semantic data integration and structured knowledge acquisition. Our Knowledge Acquisition Tool (KAT) was

built for what the article calls tribal knowledge "in Slack threads and people's heads". KAT lets subject matter experts — the actual domain experts, not just engineers — formalize their knowledge into atomic, reusable building blocks through a mind-map interface.

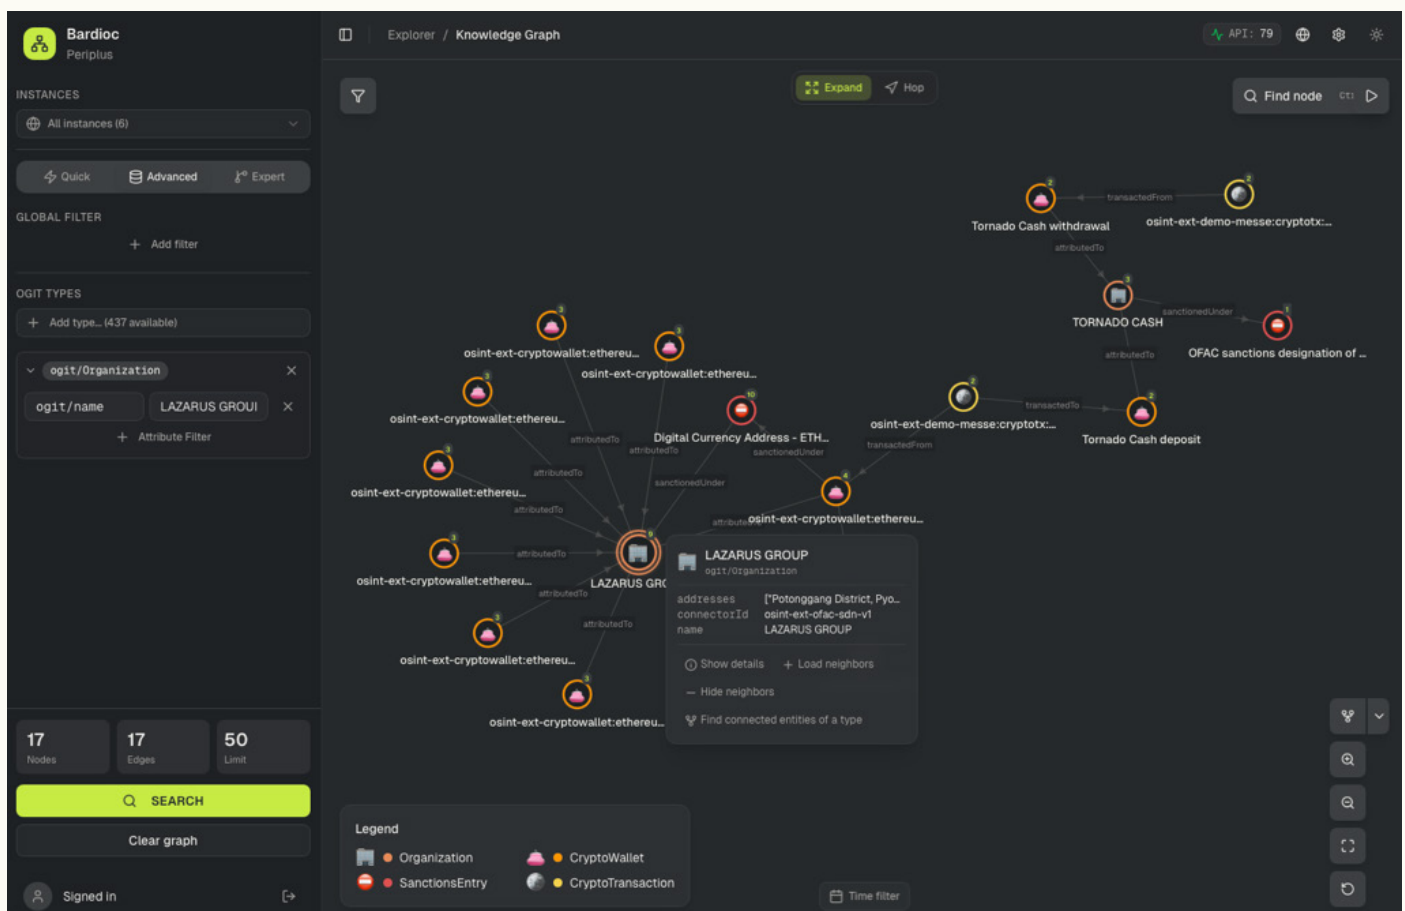
And it solves a problem that is rarely named directly: experts disagree. Two senior engineers will give two different answers to the same situation, and both can be right in their own context. Classical knowledge bases treat

that disagreement as a defect to eliminate. We treat it as the actual signal of where the hard cases live, and integrate the knowledge of multiple experts into one base — with the Reasoning Engine resolving conflicts at runtime against the goal at hand.

KAT does not ask experts to enter rules into a form. That is not how knowledge transfer works in 2026. The current research on LLM-assisted knowledge elicitation — Springer’s [Agent-in-the-Loop survey](#), the LLM-Prior framework at ICLR, work on knowledge-graph validation with humans in the loop — all point in one direction. LLMs

and human experts work hand in hand. The LLM proposes candidate rules, surfaces contradictions in the existing base, suggests ontology extensions. The human expert sharpens borders, rejects wrong patterns, adds the why behind the rule. This is where creative solutions come from. This is where root-cause explanations come from. Statistical patterns alone do not produce them. Neither does an expert sitting alone in front of a rule editor. The hand-in-hand work does. KAT is built around this dynamic, with generative AI doing the lifting on candidate generation and the human expert remaining the authority on what enters the base.


- The result is not an emergent trace but a deliberate, growing body of machine-executable organizational intelligence.



Bardioc’s Graph Explorer: navigating semantically connected data across system boundaries in real time.

Another is making compounding measurable. Foundation Capital correctly identifies the feedback loop as the central mechanism: captured decisions become searchable precedent, and every new decision adds to the graph. We made this loop explicit. Our Knowledge Impact visualization shows, in concrete terms, the exponential effect of adding knowledge building blocks to the system. The mechanism is simple. Each new building block can fire on combinations of all existing ones, not just on isolated cases. Multi-hop

reasoning over a graph is combinatorial by construction. One new piece of knowledge does not handle one more case — it compounds with existing knowledge and unlocks previously unreachable scenarios. We can show our customers exactly how each investment in knowledge transfer multiplies their automation potential. The abstract promise of compounding becomes a measurable, visual fact.



The Knowledge Impact visualization shows how each new knowledge building block compounds with existing knowledge to exponentially expand automation potential.

And there is the question of ontology as infrastructure, not afterthought. The startups Foundation Capital highlights build context graphs within specific domains — sales, finance, production engineering. Each will develop its own schema for representing entities and relationships. This is the data-silos problem all over again, one level up. OGIT is our answer: a universal, open, evolving semantic framework that provides shared meaning across domains. When a “customer” in your CRM, a “debtor” in your ERP, and a “partner” in your supply chain system all resolve to the same semantic entity in the graph, the context graph becomes truly cross-functional — not just cross-system inside one domain.

What We Learned the Hard Way

Building a production context graph at enterprise scale is a different exercise than prototyping one. A few things we learned the slow way.

The first one is bidirectional sync.

Everybody talks about “integrating data”. The real engineering work starts when the source systems are still in production use — and they always are. Changes in the graph have to flow back to the source. Conflicts have to be resolved systematically, by code, not ad-hoc. We built an entire SDK for connector development with standard conflict-resolution procedures around exactly this point. This is where most integration projects quietly stall around month four. Industry data backs this up. Gartner’s late-2025 numbers show over 50% of GenAI projects abandoned after the proof-of-concept stage. 95% of IT leaders cite integration hurdles as the top blocker. The average time from prototype to production is eight months. Integration is the wall most projects hit. We saw it play out customer by customer before we accepted that the SDK had to come first.

The second came back to OGIT being open-source.

A proprietary ontology is a worse lock-in than a proprietary database, because it locks in meaning itself. If the schema that defines what your data means is controlled by a single vendor, you have traded one silo for another, on a deeper layer. Customers and partners have to be able to inspect, extend, and contribute to the ontology. We did not make OGIT open out of idealism — it is a structural requirement for a platform that wants to represent the world over time.

The third has to do with how knowledge actually enters the system.

In a moment when LLMs feel like they can figure out anything, it is tempting to assume the AI will eventually pick up enterprise context from the data on its own. That is not what we see in production. The valuable knowledge — the exceptions, the judgment calls, the domain expertise the Foundation Capital article correctly puts “in people’s heads” — still has to be transferred deliberately. By humans who know why a particular rule exists. With LLMs working alongside them, not replacing them. What we are doing with Bardioc is bringing different AI approaches together — knowledge-based reasoning, machine learning, generative models — and using each one for the part of the problem it

actually solves. KAT is built on this assumption, and every deployment we have run so far has confirmed it.

05 An Invitation

Foundation Capital has articulated a vision we recognize. We have been building toward it from a different starting point and along a different path, and we arrived at remarkably similar conclusions about what enterprise technology needs to become.

We would welcome a conversation with Jaya Gupta, Ashu Garg, and anyone else who is thinking about context graphs, decision lineage, and semantic infrastructure. The problems are too large for any single perspective to be sufficient, and comparing notes between those who theorize and those who build is usually productive.

If you want to see Bardioc’s context graph in action — how decision paths are traced, how knowledge compounds, how a universal ontology holds it together — reach out. We are happy to show rather than tell.

→ Almato builds Bardioc, the semantic data platform. [Learn more about our platform, our capabilities, and our Reasoning Engine.](#)

→ References

01. Jaya Gupta and Ashu Garg, AI's Trillion-Dollar Opportunity: Context Graphs, Foundation Capital, December 2025. <https://foundationcapital.com/ideas/context-graphs-ais-trillion-dollar-opportunity>
02. Context Graph vs. Ontology: Differences, Roles & Use Cases, Atlan, 2026 — <https://atlan.com/know/context-graph-vs-ontology/>
03. Bijit Ghosh, Why Ontology, Context Graphs, and Decision Traces Are the New AI Substrate, Medium, 2026 — <https://medium.com/@bijit211987/why-ontology-context-graphs-and-decision-traces-are-the-new-ai-substrate-bc85e45c1ba7>
04. Beyond Context Graphs: How Ontology, Semantics, and Knowledge Graphs Define Context, Year of the Graph newsletter, Spring 2026.
05. AI Projects in I&O Stall Ahead of Meaningful ROI Returns, Gartner, April 2026.
06. Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, Gartner, June 2025.
07. Connectivity Benchmark Report 2025, MuleSoft (Salesforce).
08. Agent-in-the-loop to distill expert knowledge into AI models: a survey, Springer Artificial Intelligence Review, 2025 — <https://link.springer.com/article/10.1007/s10462-025-11255-1>
09. LLM-Prior: A Framework for Knowledge-Driven Prior Elicitation and Aggregation, ICLR 2025.
10. Knowledge graph validation by integrating LLMs and human-in-the-loop, ScienceDirect, 2025.
11. V. Voss, L. Nechepurenko, R. Schaefer, S. Bauer, Playing a Strategy Game with Knowledge-Based Reinforcement Learning, SN Computer Science 1:78 (2020) — <https://doi.org/10.1007/s42979-020-0087-8>
12. OGIT (Open Graph of IT) — <https://github.com/almatoai/OGIT>

Almato

Almato is Europe's leading specialist in AI-driven semantic data platforms. With Bardioc, Almato provides a sovereign, high-performance, AI-powered semantic data platform for the intelligent integration, analysis, and utilization of data. Based on an open ontology, the platform operates securely and transparently in line with European standards and can be deployed in various operating models, including as a Platform as a Service (PaaS) offering. It enables organizations to utilize their data flexibly and on a large scale across various operating environments.

Around 200 Almato employees develop and maintain the platform, supporting organizations in industry, the public sector, and defense in

using their data efficiently and securely, from integration and semantic modeling to AI-driven analysis and automation.

Almato collaborates with leading partners specializing in implementation and consulting, particularly in the fields of data science, security, and solution development. Together, they help customers build comprehensive, data-driven strategies.

In addition to its headquarters in Stuttgart, Almato has offices in Heidelberg, Reutlingen and Barcelona. In 2025, Bardioc was named 'AI-Powered Semantic Data Platform of the Year in Europe' by CIO Review magazine.