

Company Knowledge, Reliable Enough to Act On

Oct 6, 2026

The Zorka Lab white paper: why companies keep paying for what they already know, and how one governed model of the company changes that

Andrei Kuryan, CEO, and Stach Miklashevich, Chief Commercial Officer · Zorka Lab, Warsaw

Abstract

Every company that delivers complex work builds up a record of what it has done: the problems it solved, how it solved them, what changed and what proves it. Almost none of that record can be used. It sits in delivery reports, proposals, case studies and the memories of a few experienced people, so every function rebuilds it on its own, and now AI assistants do too.

Zorka Lab builds, verifies and maintains one machine-readable model of what a company knows, has done and can prove, so its people and its AI can act on it. The model is a business knowledge graph, built from the company's own documents and owned by the company. A governance layer, the Graph Steward, ties every claim to its source and its status. Applications run on top, and each new one reuses the same graph. The first, sales, is live. Evidence and claims assurance is available today as the Evidence Report.

1. Companies pay again and again for what they already know

Take an engineering firm of a few hundred people. Two years ago it built a predictive-maintenance system for a packaging manufacturer, and it worked: unplanned downtime fell by about a third over six months, as the client's own maintenance log shows. This week, three people in the firm need what that project taught it.

An account executive has an enquiry from a food manufacturer whose lines keep stopping. The case study on the website says the firm reduced downtime with AI-powered predictive maintenance. It doesn't say by how much or under what conditions, and the architect who knows is away until Thursday. A bid manager needs a measured result for a tender and finds a slide claiming downtime fell by half. Nobody can say where that figure came from. A project lead learns in week three that a plant network can't carry raw sensor data to the cloud. The packaging team solved exactly that two years ago, but nobody told him, because nobody knew to ask.

The firm and its clients are a composite, built to show the mechanics. They are not Zorka Lab clients, and we return to them throughout.

The company owns the knowledge; it just can't reach it. So every function pays to rebuild it.

Function	Where the cost appears
Management	Packs built on inconsistent figures and definitions
Delivery	Problems already solved get solved again
Compliance and audit	Evidence gathered by hand for every audit
HR and onboarding	Knowledge leaves with the people who hold it
Procurement	Suppliers re-assessed from scratch
Bids and proposals	Proposals without proof of what was delivered
Sales and marketing	Claims that contradict each other or the record
AI programmes	Every assistant rebuilds its context, and people re-check its answers

Knowledge workers lose about 25% of their time just searching for answers [1]. That comes to €8,500–15,700 per knowledge worker a year, spent searching for, waiting for and recreating what the company already knows [1, 3, 4]. The low end is 5.3 hours a week, the high end a quarter of the working year, both at €34.9 an hour. And 92% of organisations do not consistently capture what employees know before they retire [2].

2. Why now

AI makes answers cheap. It does not make company knowledge trustworthy. Ask a language model the bid manager's question and it may well answer that downtime fell by half, because a slide says so. An error that used to sit in one deck now appears in every answer the assistant gives.

Figure	What it says	Source
60%	AI projects without AI-ready data expected to be abandoned through 2026	Gartner, 2025 [6]
42%	Companies that abandoned most of their AI initiatives in 2025, up from 17% a year earlier	S&P Global, 2025 [7]

Figure	What it says	Source
6% vs 89%	Executives sure they can show organisation-wide AI return, against those who say AI speeds up work	Atlassian, 2026 [8]

The same research points to the fix. In one benchmark, GPT-4 answered enterprise questions correctly 16.7% of the time working directly on a database, and 54.2% of the time over a knowledge graph of the same data [5]. A 2026 study found that a semantic layer added 17 to 23 points of accuracy across three frontier models [9]. Gartner forecasts that by 2027, organisations that prioritise semantics will raise the accuracy of their AI agents by up to 80% and cut costs by up to 60% [10], and that by 2030 universal semantic layers will be treated as critical infrastructure [11]. These are single studies and forecasts, but they point the same way. Trusted knowledge is the scarce layer, and structure is how a company builds it.

3. The idea

Build company knowledge once. Keep it true. Reuse it in every workflow.

Every department keeps its own version of what the company knows, and the versions drift apart. Zorka gives the company one governed model of what it knows, has done and can prove. We think of it as the operating system for enterprise knowledge. It doesn't replace the systems of record: a CRM records that a deal closed, and a project tool records what was delivered. The systems keep the records. Zorka keeps what they mean together, with the evidence for each claim.

The model is built from documents the company has already written, not from interviews, because people forget steps and misremember causes and documents don't. Every claim carries a source, a date and a status, so the model can say "we don't know". And every piece of work done on it makes it richer.

The strongest evidence for the idea came from our own work. In one client's graph, built for its sales workflow, about 90% of the statements describe the company itself. Only one in ten is specific to sales.

Only one statement in ten is specific to sales

Statements in one client graph, built for its sales workflow, by layer




Figure 1 · one client graph, anonymised · shares rounded

The graph does not know it is a sales graph. The same knowledge answers whether to bid on a tender, which proven approach fits a new project, what the company may claim in public and what its AI assistants should work from. Sales is the first application. It is not the boundary of the company.


4. From one project to reusable knowledge

Every delivered project can be described by the same chain of eight links: customer, project, delivery, solution, capability, pain, KPI and result. The graph records each link with a source, an excerpt and a date, and describes the solution after delivery, from the evidence rather than the proposal. A typical case study keeps only four of the eight.

A case study keeps four of the eight links, and none of the evidence

Highlighted: the links that survive into a typical case study. Dashed: the links it drops.

What the project records



Every link carries a source, an excerpt and a date

What the case study keeps



No conditions, no KPI or owner, and no evidence behind the result

Figure 2 · the eight links of a delivered project, and the four a case study keeps

A case study is written to persuade, so the detail is removed on purpose. A **Business Case** keeps it: the problem and its context, the solution as delivered, the result, the evidence for that result and the conditions under which it held. A case study tells the story; a Business Case keeps the knowledge. A company can publish hundreds of case studies and still own almost nothing it can reuse.

A Business Case also separates three kinds of success. Technical success means the system works. Operational success means people use it: on the packaging lines, operators ignored alerts that said "anomaly on line 2" and acted only once the alerts named the part to replace. Business success means a metric that matters moved, and the movement can be shown. Go-live proves the first. Buyers pay for the third.

Two more levels make the knowledge reusable. A **Reusable Solution Pattern** separates what was specific to one client, such as integrations, data structures, custom rules and local regulation, from what repeats: a problem pattern, an architecture, a prerequisite, a typical risk and outcome. In the example, running the models on the factory floor and needing a year of maintenance history move into the pattern. A pattern drawn from one case stays a hypothesis until a second case confirms it. An **Offering** is what the company decides to sell again on top of the pattern, with its scope, prerequisites, delivery model and proof points.

Each level serves different people. The account executive needs the offering and its proof, the bid manager needs the verified result, and the project lead needed the pattern.

5. How the platform works

Zorka is one governed knowledge graph per company, with applications on top. A new use does not mean a new build.

One governed graph per company, with applications on top

Only the top layer is specific to a business process. A new application reuses the rest.




Figure 3 · the four layers of the platform

Sources stay where they are: documents, drives, delivery files and public data are read, not migrated. **The business knowledge graph** holds what the company knows, has done and can prove. It is governed by an ontology and owned by the client. **The Graph Steward** keeps it true (section 6). **Applications** are business processes served by co-pilots, sets of AI agents that work on the graph under the Steward's rules. Only this top layer is specific to one process. The order protects the client's investment: models will keep changing, and the governed knowledge underneath should outlive them. The ontology builds on published research, including the TRIZ Ontology presented at the TRIZ Developers Summit in 2019 and 2020 [12].

Building a graph. There are no forms and no workshops. Construction agents read the documents and map them onto the ontology while people review and steer. A build takes from 1.5 delivery days for a small company with well-organised documents, and longer as the volume grows. Before anyone commits, the company gets a graph-completeness score: how much of the required structure its documents support, and what is missing.

Reasoning. The graph decides what the model needs to see; the model decides what it means. Each task is broken into small steps, each step sees only the graph context it needs, and nothing is kept until its provenance is checked. That is how Zorka minimises hallucination risk. Errors don't become impossible, but they become traceable and can be fixed once, in the graph. The core is model-agnostic: swapping a model means re-running our benchmarks, not rebuilding the product.

Fit is treated as a conclusion, not a label. Two companies can share an industry and a size and still differ in everything that decides whether a solution transfers: the business process, the architecture, the data available, regulation and organisational maturity. So the system must be able to say no. Over any 100 researched companies, at least 30% must come back Excluded or Low, each with a traceable reason, and no more than half may come back High.

Learning. Every piece of research extends the graph. When a hypothesis meets a real buyer, it moves to Verified or Rejected; today some of these changes are recorded by hand after the conversation. When a deal becomes a delivered project, its documents become a new Business Case. The graph is refreshed on a schedule today, and event-driven monitoring is next. What compounds is this loop: each graph accumulates its company's validated history, and copying the visible architecture does not reproduce it. Client data never leaves its own graph. Learning across clients is limited to shared structure, such as types of pain and KPIs, and happens only where a client's contract allows it.

6. Keeping it true

A wrong sentence in one document is a local error. Stored as trusted company knowledge, it spreads to every decision and every AI answer that relies on it. So every claim carries its source, an excerpt, a date and a status: **Hypothesized → Reported → Verified → Rejected**.

"The client told us downtime fell by a third" is Reported. "The maintenance log shows it fell by a third over six months" is Verified. The unsourced slide claiming "by half" can only enter as a hypothesis, flagged as contradicting the log. Only one of the three belongs in a tender.

People steer; nobody overwrites the evidence. Users decide where to look and what matters. They don't type facts into the evidence: a claim entered by a person has no source, so it stays a hypothesis until a document supports it. The most valuable thing the graph shows is the gap between what a company says and what its evidence proves. If anyone could mark a claim as sound, that gap would close by assertion.

The Graph Steward enforces this. It is live with baseline governance: six of its nine functions run today, and three are next.

Function	What it does	Status
Ontology compliance	Admits only what the ontology allows	Live

Function	What it does	Status
Claim status	Keeps each claim's status, and fact, inference and hypothesis apart; some status changes are still recorded by hand	Live
Provenance and confidence	Puts a source, an excerpt and a date on every claim, with how strongly it is supported	Live
Contradiction and duplicate detection	Shows conflicting statements instead of averaging them	Live
Change traceability	Proposes every change with its source, never silently	Live
Versioning and freshness	Records when each claim was true; refreshes monthly, weekly or daily	Live
Impact analysis	Shows what depends on a fact when it changes	Next
Role-based review	Routes each change to the accountable person, who approves or rejects it but cannot write evidence	Next
Event-driven monitoring	Checks for change between scheduled refreshes	Next

7. The first proof: evidence and claims assurance

A graph doesn't only store what a company knows; it also shows what the company doesn't know. Because the ontology defines what a complete Business Case contains, building a graph exposes missing baselines, unmeasured results, unsupported claims and broken links between problem and solution.

When we built the graph for one client from its own case studies and marketing material, we weren't looking for problems. The system found 5 contradictions, 14 evidence gaps and 22 open research questions. Among them was one customer quote reused across several case studies where it didn't belong. Each finding was a marketing, legal or due-diligence risk, caught before anyone relied on it.

We offer this as the **Evidence Report**: every contradiction, gap and unsupported claim in a company's own materials, each with its source. It costs €2,500, credited in full against a graph build within 90 days. No prospect is researched and no message is written, which makes it the plainest proof that Zorka is not a sales tool.

8. The first application: sales

In complex B2B, weak sales results often have less to do with the sales process than with knowledge lost between delivery and the people who sell. Sales came first for practical reasons. It has the shortest feedback loop, since a wrong inference is contradicted by a real buyer within days. It shows a return within weeks. And the buyer already has a budget for it.

The Sales application has run on the current platform since late July 2026, as three co-pilots on one graph.

Three co-pilots share one graph. Your own tools send the message.

The Sales application, from delivered work to a validated opportunity




Figure 4 · the Sales application: three co-pilots on one graph

Case-to-Market turns delivered work into Business Cases, patterns, offerings and an ideal customer profile. **LeadGen** researches accounts and the people in them, matches each account to an offering on seven dimensions including exclusion risk, and drafts the first message. **Sales** supports discovery and records what buyers confirm or reject. For every account the answer comes down to four things: which account, which pain, which KPI and which owner. Discovery questions come from the gaps between prospect and offering; for the food manufacturer, how much maintenance history it keeps and who owns unplanned downtime.

The work costs a fraction of doing it by hand. Researching an account costs €1.40 against about €20, and a buying-committee brief €4.00 against about €110 (Growth-plan prices; manual time at €40 an hour). Screening is free, because a reasoned "no" is worth having.

We do not send. Zorka prepares everything before the send and feeds the tools a sales team already uses, such as ZoomInfo, Apollo, Clay or 11x.

Back at the engineering firm, the account executive now has the packaging Business Case, the pattern with its prerequisites and one verified claim. Nobody waits for Thursday.

9. What comes next

Today proves the engine on sales. The same graph and Steward then serve the rest of the company.

Today proves the engine on sales; the same graph then serves the company

Direction, not dates. Only what is under Today is live.




Figure 5 · the roadmap: today, next, destination

None of this carries a date. A business process becomes an application only when four things hold: it is a decision made again and again from the company's own knowledge, that knowledge is already written down, getting it wrong is expensive, and it needs new agents rather than a new ontology. The graph holds no ledger or transactions, so we make no claim on accounting or financial planning.

The destination is a company in which every function, and every AI assistant, reads from one governed model of what the company knows, and in which knowledge stays when people leave.

10. Trust: ownership, data and deployment

- **You own the graph.** Every fact in it, with its source, belongs to the client, and the client can export it.
- **Your data stays in your graph.** It is not shared with other clients or used to train a shared model.
- **Where it runs.** Zorka runs as a cloud service today. The stack is designed for deployment in the client's own cloud; private, on-premise and air-gapped deployment, and an EU-sovereign model path, are on the roadmap.
- **The specifics, in writing.** Before any build, each client gets a written answer on where its data is processed, which model providers handle it, how it is stored and protected, and how it can be exported or deleted.

- **Who can change what.** Operators steer and run the work, viewers read and comment, and nobody can overwrite the evidence.

11. How Zorka compares

Others find documents, govern data or automate answers. Zorka governs what a company can prove.

Approach	Examples	Where Zorka differs
Status quo	Experts, shared drives	A model that stays when people leave
Search and knowledge management	Glean, Microsoft Copilot, Guru	Checks what the documents actually support
Data governance	Atlan, Collibra	Governs claims about the business, not data assets
Graph platforms	Neo4j, Stardog	Built and governed for you, from your documents
Ontology platforms	Palantir	Built in weeks, at mid-market prices
In-house AI over documents	Custom retrieval builds	Sourced context that every model and use reuses
Sales context and data	Octave, Clay, ZoomInfo	Starts from what the company has delivered

These describe how each category is designed and marketed, not any vendor's roadmap. Several are complementary: a search tool can sit on top of a governed graph, and the sales data tools are what Zorka feeds.

12. Limits, and what this paper does not claim

The hard problems are real: recognising the same customer or capability under different names, deciding when a pattern transfers to a new setting, keeping a stable core while each graph adapts, and keeping every claim's validity window current. Extraction is bounded by the documents. A company whose knowledge lives only in people's heads has nothing for Zorka to read, and we say so before anyone commits.

This paper does not claim that every application described is built, that every inference is correct, that a graph alone is a moat, or that one client's results are a benchmark. Our strongest claim is also the narrowest: a governed, evidence-linked model of a business

becomes more valuable as it accumulates that business's actual experience.

13. What it is worth to a company

For an illustrative company of 500 employees, 300 of them knowledge workers, we put the value at about €275,000 a year, with every line counted at half its full value.

Where the value comes from	How it is counted	Value a year
Capacity recovered	2% of each knowledge worker's week returned	€188k
Faster onboarding	45 hires a year, each productive two weeks sooner	€59k
One workflow: bids	Half of the bids the company should never have made, avoided	€28k
Base case		€275k

At Enterprise level, Zorka costs such a company about €85,000 a year, or €107,500 in the first year including the graph build net of its credit. That is a return of about 2.6 times in the first year and 3.2 times after, at under €300 per knowledge worker a year. Errors avoided, AI projects that ship and deals won are not counted.

Conclusion

A company can only act on the knowledge it can reach. When delivered experience stays scattered, every function rebuilds it, and AI repeats whatever the documents say, contradictions included. When that experience becomes connected Business Cases, patterns, offerings and evidence, every delivered project makes the next decision easier.

It also gives management a new asset to run, and new questions to ask. How many of our significant projects have become Business Cases? Which of the capabilities we claim are backed by real deliveries? How much of what we know lives only in a few people's heads? How often do engineers stop work to answer "have we done this before"?

Sales is where Zorka starts, because it pays for the graph fastest and tests it hardest. The graph belongs to the whole company. The engineering firm in section 1 already had the answers its three people needed. It just couldn't reach them.

How to start

1. **A five-minute fit check.** We tell you whether enough is written down to build on.
2. **A graph-completeness score.** What your documents support, and what is missing.
3. **The Evidence Report.** €2,500, credited in full against a graph build within 90 days.
4. **The graph build.** Owned by you, from €1,500 to €45,000 or more, paid once, with half credited back over twelve months.
5. **Applications on the same graph.** Sales from €149 a month on self-serve plans, and from €1,200 a month by contract.

The free plan runs on public data. Operators are included in every plan, and viewers are free. Prices exclude VAT.

To talk to us: Stach Miklashevich, Chief Commercial Officer ·

stach.miklashevich@zorkalab.com

Who we are

Zorka Lab was founded by five people from semantic technology, ontology research and enterprise software. Four of them have worked together for seven years, and the team has spent two years building Zorka directly.

Founder	Role	Background
Andrei Kuryan	CEO	Ontology engineering; co-initiator of the TRIZ Ontology project; 20 years in TRIZ; formerly at PwC
Vasili Paulik	Chief Architect	25 years in commercial semantic technology at Invention Machine, IHS Markit and S&P Global
Stach Miklashevich	Chief Commercial Officer	20 years in B2B technology sales; co-founded Urb.ax and grew it from zero to \$5.5M in annual recurring revenue in under two years
Astap Kuryan	CTO	15 years in IT; Lead Software Engineer at EPAM; enterprise architecture
Sergey Boyko	CPO	Senior Director at EPAM Systems; enterprise software architecture

The platform's ontology is grounded in published research: the TRIZ Ontology, presented at the TRIZ Developers Summit in 2019 and 2020 [12].

Zorka Lab Sp. z o.o. is based in Warsaw. The current platform has been live since late July 2026.

Sources

1. [Atlassian, State of Teams 2025](#). Survey of 12,000 knowledge workers and 200 executives: leaders and teams lose 25% of their time searching for answers.
2. [APQC, Great Retirement study, 14 October 2025](#). Survey of 1,000 professionals, sponsored by eGain: 92% of organisations do not consistently capture knowledge from soon-to-be retirees.
3. [Panopto, Workplace Knowledge and Productivity Report, July 2018](#). With YouGov, 1,001 US workers: 5.3 hours a week spent waiting for information and help from colleagues.
4. [Eurostat, hourly labour costs in 2025, published 31 March 2026](#): €34.9 an hour, EU whole economy.
5. [Sequeda, Allemang and Jacob, A Benchmark to Understand the Role of Knowledge Graphs on Large Language Model's Accuracy for Question Answering on Enterprise SQL Databases, November 2023](#): 16.7% accuracy directly over SQL, 54.2% over a knowledge graph.
6. [Gartner, Lack of AI-Ready Data Puts AI Projects at Risk, 26 February 2025](#).
7. [S&P Global Market Intelligence, March 2025, as reported by Cybersecurity Dive](#): the share of businesses scrapping most of their AI initiatives rose to 42%, from 17% a year earlier.
8. [Atlassian, State of Teams 2026](#). Survey of 12,035 knowledge workers and 173 Fortune 1000 executives, January and February 2026.
9. [Rumiantsau and Fokeev, Semantic Layers for Reliable LLM-Powered Data Analytics, April 2026](#): accuracy up 17 to 23 percentage points across three frontier models.
10. [Gartner, Lack of Semantics Causes Inaccurate AI Agents and Wasted Spending, 11 May 2026](#). Forecast.
11. [Gartner, top predictions for data and analytics in 2026, 12 March 2026, as published by BigDATAwire](#). Forecast.
12. TRIZ Ontology: [Kuryan, Souchkov and Kucharavy, Towards Ontology of TRIZ, TRIZ Developers Summit 2019](#); [TRIZ Ontology: Current State and Perspectives, TRIZ Developers Summit 2020](#).