ZORKA LAB

From Engineering Knowledge to Revenue

Why the Efficiency of Complex B2B Sales Depends on More Than the Sales Process

Andrei Kuryan, Stach Miklashevich

October 2026

A company may have a strong engineering team, dozens of successfully delivered projects, and years of experience solving technologically complex customer problems — yet its sales team may use only a small fraction of that accumulated knowledge.

The problem arises on the path between those who create the solution and those who must understand to whom, when, and why it can be sold again.

As a result, the company accumulates engineering experience much faster than it develops the ability to turn that experience into reusable commercial knowledge.

> *In complex B2B, low sales efficiency may result not only from weaknesses in the sales process, but also from knowledge loss as experience is transferred from engineering and delivery to sales.*

Solving this problem requires more than tools that automate the sales process; it requires a fundamental rethink of the company's knowledge-management model. In our view, under the emerging technological paradigm, a Knowledge Graph will form the foundation of such a model.

1. Sales Has Not Only a Process Problem but a Knowledge Problem

Consider a typical company that creates complex software, engineering, or technology solutions. Over several years it has delivered dozens of projects. Those projects accumulate a large body of knowledge, including:

- what problems customers faced;
- why those problems arose;
- which requirements proved critical;
- which functions the solution had to perform;
- which architecture was selected;
- which components were used;
- which constraints had to be considered;
- which implementation difficulties emerged;
- which decisions had to be changed;
- how users applied the system;
- which metrics changed;
- what business outcome was achieved.

But this knowledge is scattered across project documentation, architecture materials, Jira, Confluence, CRM, presentations, reports, correspondence and, most importantly, the experience of engineers and architects.

At the same time, the sales team receives a significantly compressed version of that knowledge.

Client → Problem → Our Solution → Result

For selling a complex solution, this level of knowledge is critically insufficient.

This is where the gap emerges between the experience available in the company and its ability to use that experience in sales.

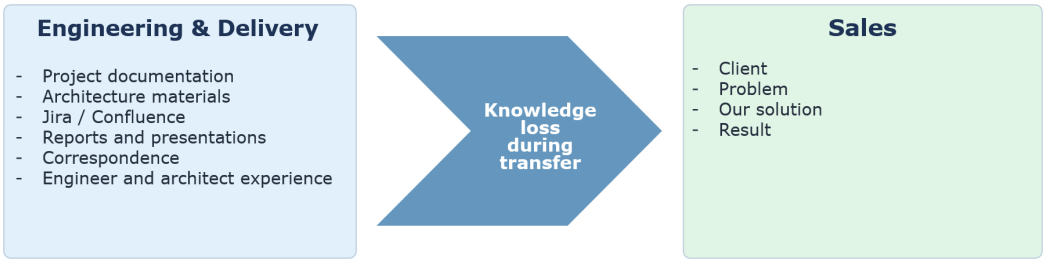

Figure 1. The Gap Between Engineering Knowledge and Sales

2. Where Engineering Knowledge Is Lost

During a real project, the company accumulates extensive knowledge and experience that often never reaches the sales team. The engineering logic can be represented as follows:

Business Problem → Business Goal → Requirements → Required Behavior → Solution Functions → Components → Implementation → Operational Use → Business Outcome → Evidence

In sales materials, this structure is often reduced to:

Client → Problem → Solution → Result

This reduction removes a significant part of the knowledge: critical facts and the relationships between them. Without those relationships, a case study remains a story about a past project rather than an instrument for the next sale.

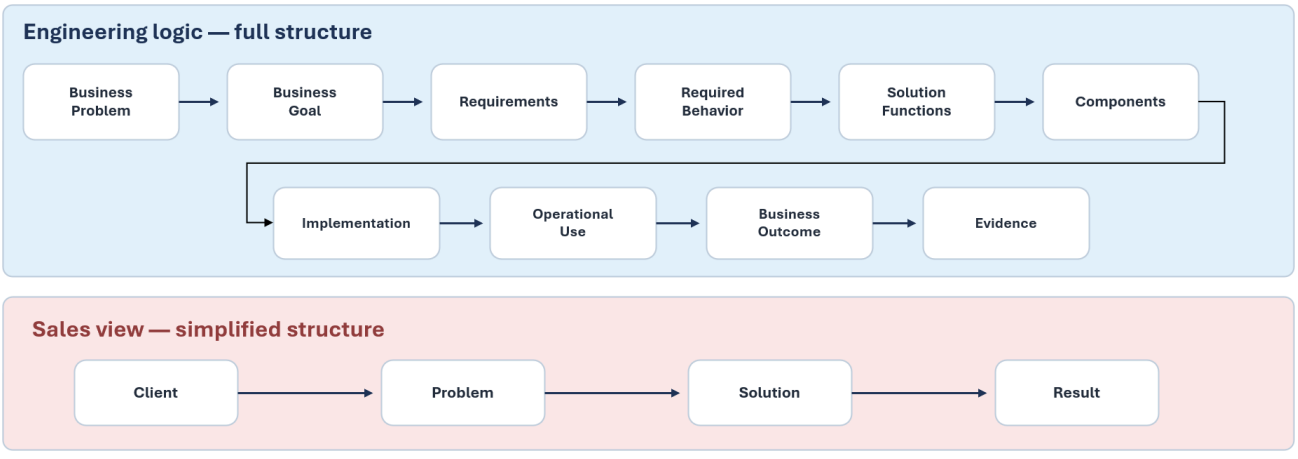

Figure 2. Full Engineering Logic vs. Simplified Sales View

3. Why Traditional Case Studies Are Not Enough

A case study is usually created as a document for a human reader. It must be concise, clear, and persuasive. As a result, much of the underlying complexity is deliberately removed.

A case study therefore answers one question well:

> *What did we do?*

But systematic commercial reuse requires answers to different questions:

- What exactly was the problem we solved?
- Why did the solution work?
- Which part of the solution can be reused?
- Under what conditions is it applicable?
- Which prerequisites are required?
- When is the solution not applicable?
- Which outcome is actually proven?
- Which elements of the case can be used as commercial claims?

That is why we distinguish between the two.

> *Case Study - a way to tell the story.*
> *Business Case - a structured knowledge model of the implemented solution, its context, applicability, and outcome.*

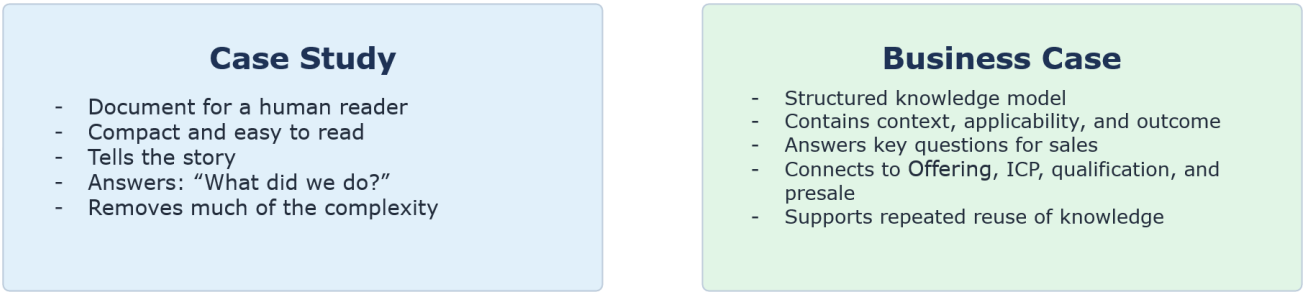

Figure 3. Case Study vs. Business Case

In Zorka Lab's approach, the Business Case links an implemented client solution to reusable commercial knowledge and then to Offering Development, ICP Development, Lead Qualification, and Presale Preparation.

This distinction is fundamental. A company may have hundreds of published case studies and still have almost no structured knowledge that can be reused systematically.

4. The Economic Consequences of Knowledge Loss

For top management, the problem becomes material when information loss turns into economic loss.

Cost of Sales Increases

The sales team must repeatedly involve solution architects, engineers, and delivery managers to determine whether the company has done something similar before, which case is relevant, whether a specific solution can be reused, and what can safely be promised to the client. Highly qualified specialists repeat work the company has already done.

The Sales Cycle Becomes Longer

The more complex the offering, the more questions require internal alignment. Instead of directly accessing accumulated knowledge, the sales team has to reconstruct it through meetings, Slack, documents, and one-to-one consultations.

Qualification Quality Declines

A conventional ICP is often described using a limited set of attributes: industry + geography + company size + technology. For a complex solution, that is not enough. Actual applicability may depend on a specific business process, problem pattern, current architecture, organizational maturity, data availability, trigger, prerequisite, constraint, or regulation. If these relationships are not represented explicitly in the ICP, the sales team qualifies companies inaccurately and superficially.

Scaling the Sales Team Does Not Scale Well

An experienced salesperson may spend years accumulating knowledge about the company's solutions. A new employee does not have that knowledge. The company becomes dependent on a small number of people who know what the company can really do and where it works. Increasing headcount therefore does not necessarily produce proportional growth in sales capacity.

The Economic Value of Completed Projects Is Lost

Every completed project should potentially increase the company's ability to sell the next projects. But that happens only when the experience becomes reusable. Otherwise, the company accumulates a history of projects without accumulating commercial knowledge at the same rate.

> *A company can successfully accumulate project experience while failing to accumulate reusable business knowledge at the same pace.*

5. Business Case as a Bridge Between Engineering and Sales

The documents and presentations that engineering passes to sales may already contain gaps or even errors in knowledge. The company needs an artifact that preserves this knowledge during transfer.

In Zorka Lab's approach, a Business Case is represented as a subgraph of the shared Knowledge Graph.

Client Context → Business Problem → Business Goal → Requirements → Solution Functions → Implemented Solution → Operational Use → Business Outcome → Evidence

This preserves the causal and semantic structure of the project.

For example, sales sees not only that the company built and deployed a system for a client, but also:

- what problem it addressed;
- which solution functions were required;
- what capability the client organization gained;
- how the system was used;
- which process changed;
- what outcome was achieved;
- how well that outcome is supported by evidence.

At Zorka Lab, when building a Business Case we also distinguish technical, operational, and business success: client acceptance of a solution, or even Go-Live, does not by itself demonstrate realized business value.

For sales, this distinction is critical. It makes it possible to separate “we implemented a function” from “this function helped produce a specific business outcome.”

6. One Business Case Is Not Enough

Even a well-structured Business Case still describes one specific client. The next task is to determine which part of that experience can be reused.

Reusable Solution Pattern

A concrete project must be separated into what was client-specific and what is repeatable.

Client-specific elements may include: integrations with other systems; a unique data structure; client-specific business rules; customization; specific regulatory requirements.

Repeatable elements may include: problem pattern; reusable function; reusable component; architecture pattern; implementation pattern; prerequisite; typical risk; outcome pattern.

Therefore, the knowledge path from engineering to sales should not look like:

Project → Sales Presentation

but rather:

Project → Business Case → Reusable Solution Pattern

In the Zorka Lab model, a Reusable Solution Pattern includes a Problem Pattern, Context Pattern, Solution Pattern, Reusable Functions, Implementation Pattern, Outcome Pattern, applicability conditions, prerequisites, exclusion conditions, and replication constraints.

This is no longer merely a description of a past project. It is a model of how the experience may be reused.

It is important to note that a pattern extracted from a single Business Case is still a hypothesis of reuse and must be validated by other cases or by market data.

7. From Reusable Pattern to Offering

The next step is commercialization.

Even when a company knows that part of a solution can be reused, it still has to decide:

> *What exactly are we prepared to offer the next client?*

This is where the Offering appears. It answers the following questions:

- which problem we undertake to solve;
- in what context;
- which capabilities we provide;
- which functions are included in the solution;
- what is included in scope;
- what is excluded;
- which prerequisites exist;
- which delivery model is used;

- which outcome is expected;
- which proof points support it.

In the Zorka Lab model, the Offering is a separate layer of the shared Knowledge Graph, derived from a Reusable Solution Pattern rather than copied from a solution implemented for a single client.

The company therefore works across three different levels of abstraction:

Business Case → Reusable Solution Pattern → Offering

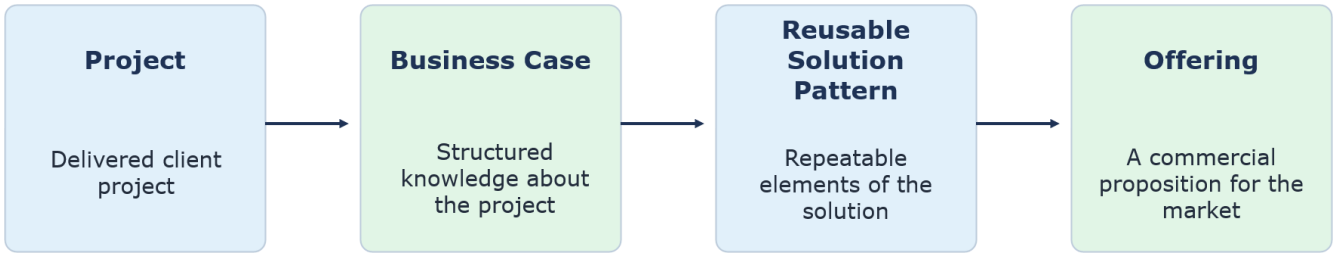

Figure 4. From Project to Commercial Offering

The Business Case answers what happened for a specific client. The Reusable Solution Pattern answers what can be repeated from that experience. The Offering answers what the company is prepared to sell repeatedly.

Confusing these three levels is one reason why complex B2B sales scale poorly.

8. Why a Knowledge Graph Is Needed

Everything described above could theoretically be stored in documents. But then the primary asset — the relationships between pieces of knowledge — would remain implicit.

A Knowledge Graph represents knowledge as a set of objects and the relationships between them. Meaning for decision-making emerges through links between a signal, problem, KPI, capability, evidence, and outcome.

For a human reader, these relationships can be described in text. For AI, automated reasoning, and reuse across different processes, they must be represented explicitly.

9. What Changes for Sales

At this point, the Knowledge Graph becomes not a repository but a decision-support layer that can answer real commercial questions.

Who Should We Target?

Instead of a superficial ICP, the company can use substantive indicators of solution applicability.

Problem Pattern + Context Pattern + Trigger + Prerequisites + Technology Environment + Exclusion Conditions

Is This Prospect a Fit for Us?

Knowledge about the prospect can be compared with the Offering to evaluate Problem Fit, Context Fit, Technology Fit, Prerequisite Fit, Persona Fit, Trigger Fit, and Exclusion Risk.

Prospect Context ↔ Offering Target Context

Which Case Should We Use?

A case does not have to be selected only from the same industry. Similarity can be evaluated by problem, business process, context, function, capability, architecture, and outcome.

What Can We Claim to the Client?

A commercial claim is linked to a proof point and evidence. Sales can distinguish a verified result from a reported claim, inference, generalization, or hypothesis.

Sales Claim → Proof Point → Evidence

What Should We Ask in Discovery?

Discovery questions can be generated not from a generic sales playbook but from specific gaps between the Prospect and the Offering: data-source availability, integration prerequisites, problem severity, KPI ownership, business trigger, and organizational capability.

As a result, sales begins to use not only product information, but structured knowledge of the company's accumulated experience.

10. A Knowledge Graph Makes Visible What the Company Does Not Know

One important function of a Knowledge Graph is not only to store knowledge but also to reveal what is missing.

If the Business Case structure is defined in advance, it becomes possible to see:

- where the baseline is unknown;
- where evidence is missing;
- where the business outcome has not been measured;
- where it is unclear which function produced the result;
- where an applicability condition is missing;
- where a prerequisite has not been defined;
- where a claim is unsupported;
- where traceability between the problem and the solution is missing.

In the Zorka Lab model, this is supported by Case Assessment, Information Gap, Business Case Issue, and dimensions of completeness. We can detect gaps because we apply an ontology that defines the expected structure and completeness rules.

> *The company can see not only what it knows, but also which critical knowledge it is missing.*

The Knowledge Graph makes this a subject of systematic management.

11. How Zorka Lab Builds the First Knowledge Graph

The Zorka Lab approach begins with building a Business Case Knowledge Graph.

Source materials may include:

- client cases;
- project documentation;
- product documentation;
- architecture materials;
- delivery artifacts;
- expert knowledge;
- existing sales materials.

They are transformed into a Knowledge Graph based on an ontology.

Ontology

The ontology defines which types of entities and relationships are required to describe a Business Case. Examples include Business Problem, Requirement, Solution Function, Capability, Business Outcome, Evidence, Reusable Pattern, and Offering.

Knowledge Graph

The Knowledge Graph contains the company's specific knowledge: concrete projects, problems, solutions, functions, outcomes, evidence, and Offerings.

> *The Ontology describes what should be present in the Knowledge Graph.*
> *The Knowledge Graph describes what the company actually knows.*

12. Building the Graph Also Becomes a Knowledge Audit

The initial generation of the Knowledge Graph is valuable in itself because it reveals the quality of the accumulated knowledge.

It can expose:

- incomplete Business Cases;
- contradictory descriptions;
- unsupported claims;
- missing metrics;
- missing evidence;
- gaps between requirements and solution;
- missing applicability descriptions;
- missing reusable elements;
- weakly described business outcomes.

> *We know what knowledge we have, how reliable it is, and what is missing.*

For a technology company, this is already a management asset in its own right.

13. Human Validation: AI Must Not Turn an Assumption into a Fact

AI can significantly accelerate the process of building a Knowledge Graph.

The model can:

- extract statements from documents;
- classify them;
- map them to ontology concepts;

- find gaps;
- detect contradictions;
- propose generalizations;
- form candidate reusable patterns.

But this creates a fundamental risk: AI may produce a reasonable inference that is not actually present in the source data.

The following categories therefore need to be distinguished explicitly:

Fact | Reported Claim | Inference | Generalization | Hypothesis

> *AI structures and proposes. Experts validate.*

An engineer validates technical facts. An architect validates the solution structure. A business owner validates business outcomes. The responsible manager validates the commercial interpretation.

This preserves the boundary between automation of knowledge work and accountability for the content of that knowledge.

14. The Knowledge Graph Must Evolve with the Company

A newly created graph will quickly become outdated. When the sales team begins selling, it receives market signals that either confirm or refute the hypotheses represented in the Knowledge Graph.

After a deal is closed and the subsequent delivery is completed, new facts emerge about the client's problem and the specifics of solving it. Those facts must also enter the Knowledge Graph. Once delivery is complete, a new commercial cycle begins, carrying forward the graph with all updates accumulated during the previous cycle.

The Knowledge Graph should therefore be treated not as a one-off data project, but as an evolving knowledge asset.

In the Zorka Lab architecture, this task is intended to be supported by the Graph Steward.

Its role is to maintain:

- ontology consistency;
- provenance;
- knowledge status;
- freshness;
- versioning;
- contradiction detection;
- change traceability;
- impact analysis;
- human validation.

> *The primary task of the Graph Steward is to keep the company's commercial knowledge current, consistent, verifiable, and reusable.*

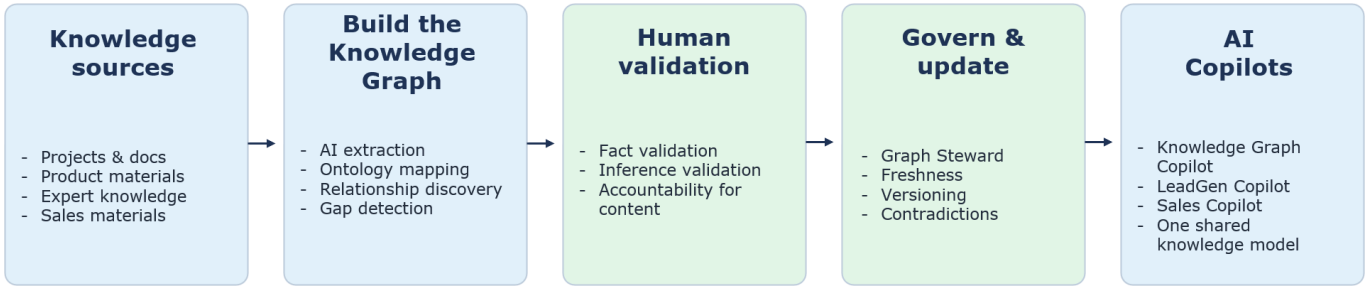

Figure 5. How Zorka Lab Builds and Maintains Knowledge Graph

15. What Changes for Top Management

This approach introduces a new object of management.

Management usually governs functional areas such as Engineering, Delivery, Product, Marketing, and Sales.

But between these functions there is another critical asset:

> *Reusable Company Knowledge*

This leads to a new set of management questions:

- what share of significant projects has been transformed into Business Cases;
- which Business Cases are complete enough for commercial use;
- which solution capabilities are supported by real implementations;
- which outcomes have evidence;
- which knowledge gaps recur;
- which Offerings are backed by proven experience;
- how much knowledge exists only in the heads of a few specialists;
- how dependent sales is on manual engineering involvement;
- whether a new salesperson can use the company's experience without constantly going back to the people who created it.

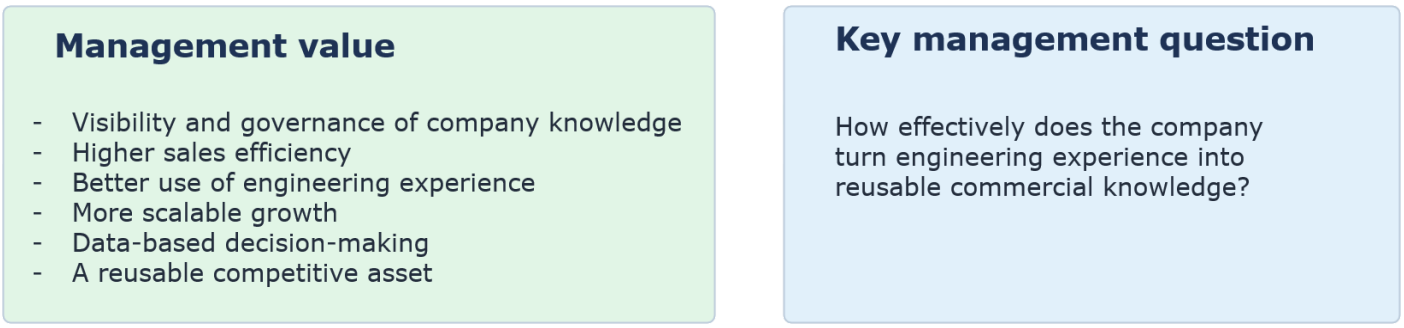

Figure 6. What This Gives Top Management

This fundamentally changes how the problem is framed.

> *How effectively does the company turn engineering experience into reusable commercial knowledge?*

And only after that: how effectively does sales use that knowledge?

16. From Engineering Experience to a Scalable Sales Capability

For companies that create complex B2B solutions, sales cannot operate independently of engineering knowledge. The more complex the solution, the stronger this dependency becomes.

Scaling sales therefore requires more than scaling the sales process. It requires scaling access to the company's accumulated knowledge.

The transformation can be represented as follows:

Engineering & Delivery → Business Case KG → Reusable Solution Patterns → Offerings → LeadGen & Sales

Engineering & Delivery create experience. Business Cases structure it. Reusable Solution Patterns isolate repeatable knowledge. Offerings define what the company can sell. The Knowledge Graph connects knowledge, evidence, and applicability. LeadGen & Sales use this knowledge to make decisions.

With this approach, the value of a completed project does not end with delivery. The project becomes a source of knowledge for subsequent sales opportunities.

Conclusion

Sales can effectively sell only what the company can make available in a form that can actually be used.

If accumulated engineering and delivery experience remains scattered across documents and people, sales must continually reconstruct it from scratch.

If that experience is transformed into connected Business Cases, Reusable Solution Patterns, Offerings, Evidence, and Sales Knowledge, the company begins to build a reusable commercial asset.

The Knowledge Graph provides the required structure. It preserves not only facts but also the relationships between them.

Problem → Requirement → Function → Solution → Usage → Outcome → Evidence → Reuse → Offering

These relationships turn the history of a completed project into knowledge that can be used in the next sale.

For a technology B2B company, a Knowledge Graph is therefore not merely a new way to organize data. It is an infrastructure for transferring knowledge from engineering to the market.

> *Zorka Lab turns the company's engineering and delivery experience into a governed Knowledge Graph that can be systematically used in Lead Generation and Sales.*

The transition from project experience to reusable commercial knowledge is what creates the foundation for scalable sales of complex B2B solutions.