

Zenskar: The AI-Native Platform Built After the LLM Revolution

A Thank You to Our Partners in Innovation

We're grateful to Zenskar for participating in the Next Generation Solution Showcase on Billing and Revenue Management. This blog is written for CFOs of fast-growth companies who are building the foundation of their finance tech stack. Leaders who need crystal-clear visibility into cash flow, want data-driven investment strategies, and require accurate forecasting to guide their next phase of growth.

"Agents That Work as an Extension to Your Team"

Zenskar describes itself as an AI-native order-to-cash platform. Apurv Bansal, Co-founder and CEO, uses a revealing analogy: "If you think of AI as the engine of a Ferrari, our architecture is the body of the Ferrari that we've built out over a couple of years."

He further stated, "Legacy billing platforms were architected in 2014, before LLMs existed. When they add "AI features" today, they're bolting GPT onto systems never designed for it."

Apurv then moved to a different question: "*If LLMs can understand context and make decisions, what becomes possible in order-to-cash?*" They intentionally architected Zenskar from the ground up to leverage LLMs, giving them what Apurv calls a "structural advantage" over legacy players built a decade ago.

The result? Agents that autonomously manage processes within order-to-cash, working like members of your finance team to help you scale without scaling headcount.

Who Zenskar Serves Best

Zenskar is purpose-built for companies operating in the AI economy, where pricing has gotten dramatically complex.

You're selling hundreds of SKUs in various bundles with matrix pricing that varies by geography, managing minimum guarantees and overage fees, and pulling usage data from your warehouse before billing, and mid-cycle contract amendments happen constantly.

Your finance team is small, maybe four to five people managing billing & revenue recognition for a \$100 million revenue company. They're spending most of their time chasing sales for clarifications, reconciling product names manually, and fixing under-billing in spreadsheets.

You're facing the accounting workforce shortage, which means you can't just throw more bodies at the problem. You need automation that works with your complexity, not systems that force you to simplify your business model.

If your order-to-cash runs on spreadsheets, homegrown tools, and manual processes that leak revenue, Zenskar was designed for you.

What Makes Zenskar Unique: Born Post-LLM Era

Zenskar has "the luxury of being born post the LLM era." While legacy billing platforms were architected a decade ago and are now trying to bolt AI onto systems never designed for it, Zenskar built its entire architecture bottom-up with LLMs in mind.

The platform is model agnostic, meaning as newer, better AI models emerge, Zenskar can leverage them without architectural rewrites. But Zenskar is thoughtful about where AI helps. Bansal is clear: "AI is not magic." For processes requiring 100% accuracy like creating journal entries and posting to your ERP, they use deterministic code. For ancillary processes where a human can review output, AI agents do the heavy lifting.

The difference shows up in three architectural advantages:

1. Contracts AI that actually understands context

The Contracts AI doesn't just use OCR to read text. It contextually understands contracts, extracting billing-relevant information while ignoring irrelevant clauses. When it encounters a 20-page MSA, it pulls out the \$2,000 flat fee from a table and the two-cent overage fee buried in paragraph 47.

2. Graphical data model for unlimited pricing complexity

Traditional platforms store pricing linearly—rows in a database with fixed prices. Zenskar uses a graphical data model that captures dependencies, hierarchies, and conditional logic natively. This means complex structures like "GPT-3.5 at \$0.002 per token, GPT-4 at \$0.006, 15% discount above \$1,000 monthly spend, credit bundles consuming at different rates by model" are configurable through the UI without custom code.

3. Decoupled metering and billing

The decoupling of contracts and usage is another architectural win. Most billing tools couple these together, requiring your engineers to embed billing logic when sending consumption data. Zenskar decouples them. Engineering connects your database once via native connectors, APIs, or CSV uploads. Finance can then aggregate and create billable metrics independently, dramatically lowering engineering effort and enabling pricing experiments without touching the data pipeline.

The Capability That Changes Everything: Custom-Prompted Agents

Zenskar's agents can be trained with custom prompts specific to your business.

The Insights Agent lets you talk to your data conversationally. Ask "*What's my MRR?*" and it converts natural language into structured queries and visualizes results. But the power comes when you teach it your business rules. Tell it "When calculating ARR, exclude implementation fees" and it learns.

The Action Agent takes commands from plain English. You talk to it like an analyst: "Auto-remind enterprise customers who are 15 days past due" or "Issue credit notes in bulk to customers affected by the pricing error," and it executes.

Custom prompts solve problems other platforms can't touch. When product names aren't standardized across sales contracts, you can train the agent to map "Voice API," "Voice Calling API," and "Voice Minutes" to the single standardized name in your ERP. (More on this in the customer story below.)

Customer Story: Eliminating the Weekly Finance-Sales Bottleneck

The Situation

A conversational AI platform serving financial institutions faced a problem killing their finance team's productivity. They had multiple products serving a few hundred customers with five people managing order-to-cash.

The core issue? Product name standardization. Their sales team used different names for identical products across contracts. The Voice API product appeared as "Voice API" in some contracts, "Voice Calling API" in others, and "Voice Minutes" in yet others. The ERP had a single structured product name.

Every month, finance would receive contracts and couldn't reconcile product names with the ERP. This triggered a week-long back-and-forth with sales for clarification. The lack of standardization delayed billing, revenue recognition, cash collection, and book closing.

The Action

The company implemented Zenskar's Contracts Agent with custom prompts. They trained the agent: "If you encounter names like 'Voice Calling API' or 'Voice Minutes' across sales contracts, map them into the Voice API product in Zenskar, since it's the standardized name used in ERP."

The agent would read through each PDF contract. When it encountered any of the non-standard product names, it automatically mapped them to the correct ERP product name.

The Results and Why They Matter

The weekly back-and-forth between finance and sales disappeared entirely. That's four to five days of productivity returned every single month. At a fully-loaded cost of **\$120K** for a finance manager, that's roughly **\$25K in annual value** from eliminated waste.

Invoicing became automated because there was no longer ambiguity about which product to bill. The under-billing that occurred from manual errors stopped. The company eliminated this entire category of revenue leakage.

But the strategic impact went deeper. The finance leader could finally spend time on activities that moved the business forward. Instead of playing detective with product names, they could analyze which products had the best retention and which customer segments were most profitable.

The Bottom Line: Architecture Matters

For CFOs evaluating billing platforms, Zenskar represents a fundamentally different approach. This isn't a legacy system trying to add AI features. This is a platform architected from day one to leverage AI as a core capability.

The ability to train agents with custom prompts means the platform adapts to your business instead of forcing you to adapt. The model-agnostic architecture means as AI improves, Zenskar gets better automatically.

Watch their full Next Generation Solution Showcase presentation to see the agents in action. Visit zenskar.com to learn more.

Your finance team shouldn't be spending 80% of their time on operational grunt work. They should be driving strategic decisions that accelerate growth.