

EBOOK

4 GenAI Use Cases That Power Intelligent, Scalable, and Context-Aware Systems

How a Unified Database Foundation Is Your
Competitive Edge in the Agentic AI Era

Dify **TOPRISM**  **APTSell**

Introduction


Why GenAI Demands a New Database Foundation

Generative AI (GenAI) is forcing a fundamental shift in application architecture. While large language models (LLMs) have captured the world's imagination, the real challenge has emerged from behind the scenes: the data infrastructure. An AI application is only as intelligent as the fresh, contextual data it can access in real-time.

In the race to innovate, many have created a “Frankenstein stack” of separate databases: an operational database for the application, a vector database for semantic search, and often a caching layer to stitch them together. This complexity creates significant roadblocks:

- **Data Freshness Issues:** Constant, failure-prone ETL pipelines are needed to sync data, leading to stale context and less accurate AI responses.
- **Architectural Complexity:** Managing multiple systems increases operational overhead, slows down development cycles, and inflates the total cost of ownership (TCO).
- **Limited Queries:** It's nearly impossible to run a single query that combines real-time transactional data with semantic vector search, limiting the intelligence of your application.

To build truly AI-native systems, you need a new foundation. This eBook showcases four real-world GenAI use cases from companies that have embraced a unified approach with TiDB. You'll learn how simplifying the data stack is the key to unlocking the full potential of GenAI.



1

Developer Productivity

Development teams are under immense pressure to ship intelligent features faster than ever. However, they are often slowed down by infrastructure complexity. A fragmented data stack—with separate databases for transactions, vectors, and caching—creates significant cognitive overhead. This requires fragile data pipelines, and forces engineers to spend more time managing infrastructure than building the product.

A unified data platform is the key to unlocking developer productivity. By simplifying the architecture, it eliminates data silos and complex sync jobs, allowing developers to focus on what they do best: creating innovative, AI-powered user experiences.

Dify.AI: Scaling an Open Platform for Agentic AI Builders

As their platform's adoption exploded, the Dify.AI team found themselves managing a growingly complex infrastructure of eight separate database containers. This fragmented architecture was a major bottleneck, creating data silos and complex synchronization jobs that hampered both scalability and developer productivity.

The TiDB Advantage

Dify.AI migrated its entire backend to a single, unified TiDB cluster. This strategic move consolidated all application metadata, user context, and vector embeddings into one system. With TiDB's seamless horizontal scalability and strong consistency, Dify.AI can now support its growing user base without worrying about database performance.

Outcomes

- **Unified Architecture:** A single source of truth for all structured and semantic data.
- **Seamless Scalability:** Effortless horizontal scaling to meet growing demand.
- **Simplified Stack:** Accelerated performance and reduced operational complexity for developers.

“TiDB allowed us to consolidate our entire backend, letting our engineers focus on building AI features instead of managing database complexity. It was a game-changer for our productivity.”

— Head of Platform, Dify.AI

Dify

[READ THE FULL CASE STUDY »](#)



2

BI & Analytics

Business Intelligence is moving beyond static dashboards. The new frontier is conversational analytics, where any user can ask complex questions in natural language and get immediate, data-backed answers. Delivering this experience requires solving a difficult technical problem: performing a hybrid workload that combines fast, transactional queries with real-time semantic search.

Legacy systems struggle with this, forcing developers to build complicated workarounds. A context-aware data platform must handle these hybrid queries natively to provide the speed and accuracy needed for a fluid, interactive user experience.

TOPRISM: Delivering Enhanced Business Intelligence with Chat2Query

To accurately translate a question like, “Which products sold best in the northeast last quarter?” into a precise SQL query, the system must process a difficult hybrid workload: combining fast transactional processing with real-time vector retrieval for semantic understanding.

The TiDB Advantage

Toprism leverages TiDB’s unified data engine, which combines vector search with its core SQL capabilities. This allows them to run a single, powerful query that uses vector similarity to understand user intent and structured SQL filters to retrieve the exact data needed, all at once.

Outcomes

- **Reduced Latency:** Faster, more accurate responses for complex Chat2Query requests.
- **Simplified Infrastructure:** A single cloud-native database replaced a multi-system stack.
- **Accelerated Deployment:** Faster AI development cycles with minimal DevOps overhead.

“To truly democratize data access, you need both semantic understanding and SQL-level precision. TiDB’s ability to do both in a single query is what makes real-time conversational analytics possible.”

— CTO, TOPRISM

TOPRISM

[READ THE FULL CASE STUDY »](#)



3

Sales Enablement

Today's sales teams are drowning in data but starving for insights. Critical information is often fragmented across data silos: the CRM, marketing automation platforms, internal knowledge bases, and spreadsheets. This prevents sales professionals from getting a 360-degree view of the customer, and forces them into inefficient, manual workflows.

A truly intelligent sales enablement platform must do more than just present data; it must unify information and automate workflows to proactively assist the sales team, turning raw data into actionable intelligence and guided selling.

APTSell: Building an AI Digital Sales Employee

APTSell's system needed to do more than just answer questions; it had to proactively assist sales teams by orchestrating complex, context-aware workflows. This required unifying diverse data (i.e., CRM, docs, analytics) and intelligently acting on it.

The TiDB Advantage

APTSell built its platform using TiDB Starter as the unified data core (the AI's "brain") and TiDB AutoFlow as the intelligent "nervous system." AutoFlow leverages the unified SQL, Vector, and Graph data in TiDB to dynamically generate and execute optimal sales task flows automatically.

Outcomes

- **Intelligent Automation:** Automated, context-aware workflows significantly boost sales productivity.
- **Unified Knowledge:** A real-time knowledge graph provides instant, accurate answers.
- **Centralized Data:** Real-time analytics on unified data enable hyper-personalized engagement.

“We didn't want to just build another dashboard. With TiDB's unified data and workflow automation, we were able to create a truly intelligent digital employee that proactively helps our users win.”

— CEO, APTSell



[READ THE FULL CASE STUDY »](#)

4

Retail Operations

For large retailers, access to real-time information across thousands of distributed stores and operational centers is a massive challenge. When in-store staff or IT support need immediate answers to procedural or technical questions, delays and inaccuracies lead to poor customer service and high operational costs.

Standard knowledge retrieval systems often fail because they lack context. They can find documents with keywords but can't understand the critical relationships between products, store locations, and known issues. This makes them unreliable for complex, mission-critical problem-solving.

New Retail Giant: Transforming Knowledge Access with GraphRAG

Standard Retrieval-Augmented Generation (RAG) was not enough for this new retail giant. It could find documents with keywords, but it couldn't understand the critical relationships between data points. Ultimately, it failed to deliver the precise, context-aware answers needed for complex operational issues.

The TiDB Advantage

This retailer implemented a cutting-edge GraphRAG architecture using TiDB. By storing its knowledge base as both a structured graph and vector embeddings within the database, its AI assistant gained a deeper level of contextual understanding, using the graph to retrieve more precise semantic context for the LLM.

Outcomes

- **70% reduction** in IT support ticket resolution time.
- **45% cost savings** by consolidating documentation platforms.
- **40% reduction in training** time for employees via AI-enhanced SOP access.

“With GraphRAG on TiDB, we moved from simple keyword search to true contextual understanding. The impact on our support efficiency and operational costs was immediate and profound.”

— VP of IT Operations

[READ THE FULL CASE STUDY »](#)

Conclusion

Scaling Intelligence Through Unified Data Infrastructure

Across developer tools, business intelligence, sales, and retail, a clear pattern emerges. The success of any serious GenAI application depends on fast, reliable, and scalable access to unified, real-time context.

Trying to build the future on the fragmented data stack of the past is a losing battle. As these four companies demonstrate, a unified data foundation is the solution. It simplifies architecture, accelerates development, and unlocks the kind of transformative user experiences that will define the next generation of software.

Ready to build your GenAI future?

Your most ambitious AI initiatives require a data foundation built for the challenge. Let our experts show you how a unified, real-time database can help you simplify your stack and accelerate your roadmap.

Schedule a Personalized Tech Talk



TiDB, powered by PingCAP, unlocks limitless scale for data-intensive businesses. Our advanced distributed SQL database enables leading enterprises, SaaS, and digital native companies to build petabyte-grade clusters while managing millions of tables, concurrent connections, frequent schema changes, and zero-downtime scaling. Large organizations who have adopted TiDB, such as Databricks, Pinterest, and Plaid, are finally free to focus on their future growth instead of complex data infrastructure management. With AI-driven innovations and multi-cloud flexibility, TiDB offers unmatched agility, resilience, and security.

To learn more, please visit:

[www.TiDB.io](https://www.tidb.io)

