TiDB

WHITE PAPER

Architecting FinTech Resilience

How Distributed SQL Unlocks Scalable Innovation with Zero Downtime

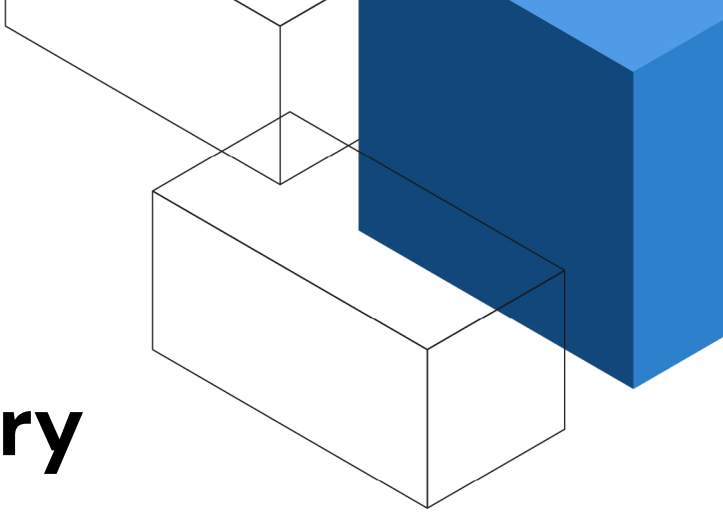

Executive Summary

FinTech innovation is redefining the speed, intelligence, and personalization of financial services—but legacy data architectures are holding teams back. As customer expectations rise and regulatory environments grow more complex, engineering leaders need infrastructure that scales elastically, operates reliably, and supports real-time, AI-enabled applications.

This white paper explores why distributed SQL has become a foundational technology for next-generation FinTech platforms—and how TiDB, an open-source, cloud-native distributed SQL database, is enabling this shift. Designed for horizontal scale, high availability, and mixed workload processing, TiDB offers an alternative to monolithic RDBMSs and brittle sharded systems.

Key themes covered include:

- The scalability, resiliency, and multi-tenant requirements of FinTech platforms
- Why traditional databases struggle with cross-region deployments, real-time workloads, and schema agility
- How distributed SQL combines strong consistency with horizontal scale and SQL compatibility
- TiDB's architecture advantages, including built-in fault tolerance, elastic operations, and integrated analytics
- Success stories from FinTech leaders like Plaid, Square, and SB Payment Service who use TiDB to accelerate innovation

Whether you're scaling a digital bank, building an embedded finance platform, or enabling intelligent automation in compliance or payments, this paper outlines the architectural playbook for resilient, future-ready FinTech systems.

Contents

Introduction

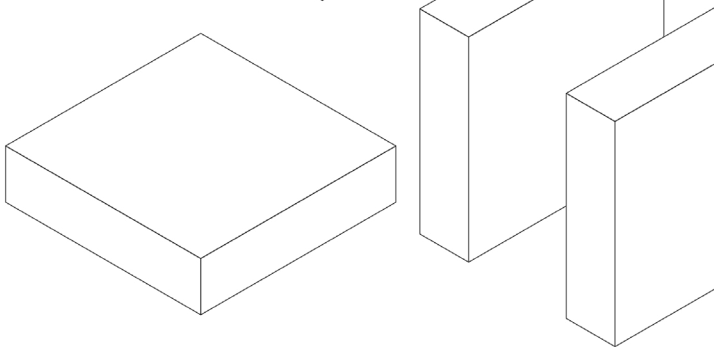

The global FinTech market, valued at $340.10 billion in 2024, is projected to soar to $1,126.64 billion by 2032, expanding at a Compound Annual Growth Rate (CAGR) of 16.2%.[1] This trajectory underscores the immense scale of data generation and the critical need for robust data infrastructure. Traditional database systems, conceived for a less dynamic era, are increasingly proving inadequate, leading to performance bottlenecks, scalability limitations, and escalating operational costs.

In this demanding environment, TiDB emerges as a transformative solution. TiDB is an open-source, distributed SQL database meticulously engineered to bridge the critical gap between the horizontal scalability characteristic of NoSQL systems and the transactional reliability (ACID compliance) of traditional relational databases. By offering a unique combination of horizontal scalability, strong data consistency ensured through the Raft consensus algorithm, and powerful mixed workload processing capabilities, TiDB empowers FinTech organizations to overcome their most pressing technical hurdles.

The velocity of FinTech market growth signifies that data infrastructure is no longer a mere supporting element but a core competitive differentiator. A unified platform approach, such as that offered by TiDB, becomes essential for navigating the complexities of high data volumes, real-time processing needs, and rigorous regulatory scrutiny

This white paper explores how distributed SQL—specifically TiDB—solves these challenges by delivering a cloud-native, highly available, and horizontally scalable data foundation for FinTech innovation. Readers will learn:

- Why FinTech platforms are outgrowing monolithic and vertically scaled database systems
- How distributed SQL enables multi-tenant scalability, strong consistency, and real-time responsiveness
- What distinguishes TiDB from legacy RDBMS and NoSQL alternatives
- How features like mixed workload processing, MySQL compatibility, and horizontal elasticity support evolving regulatory and performance needs
- Real-world success stories from FinTech innovators who've deployed TiDB to streamline operations, accelerate development, and scale securely

Whether you're a CTO rethinking platform design, a data architect modernizing backend systems, or an engineering leader scaling AI-native applications, this guide will provide practical insights into building resilient infrastructure for financial innovation.

[1] 35 Fintech Statistics: Market Growth, Investment, and Trends, April 14, 2025, https://www.docuclipper.com/blog/fintech-statistics/.

CHAPTER 1

The Five Growth Challenges Facing FinTech Companies

As the FinTech industry continues its meteoric rise—projected to surpass $1.1 trillion in global market value by 2032—the pressure on data infrastructure has never been greater. Rapid digitalization, embedded finance, and AI-driven services are transforming financial products, but they're also pushing backend systems to their limits.

For engineering teams, this means navigating far more than feature velocity. Today's platforms must ingest and process millions of transactions in real time, deliver sub-second insights, maintain data integrity across distributed systems, and comply with increasingly strict global regulations. Legacy database architectures, designed for simpler times, weren't built for this scale, complexity, or risk profile.

This chapter outlines the five most urgent challenges facing FinTech platforms—and why solving them requires more than incremental fixes.

Managing High-Volume Throughput

FinTech platforms must process millions of transactions per second with low latency. The digital payments segment is expanding at a CAGR of 19.4%, with instant payment values projected to reach $58 trillion by 2028.[1] The U.S. ACH Network processed 33.6 billion payments in 2024, totaling $86.2 trillion.[1]

Legacy databases struggle under such demands, becoming performance bottlenecks. This is acute during peak loads. One high-traffic payment FinTech faced limitations with its on-premise servers during national events due to infrastructure inflexibility. The "millions of transactions per second" benchmark is continuously escalating, necessitating a database architecture capable of predictable and cost-effective scaling.[1] System outages or slowdowns can severely impact brand reputation and customer loyalty.

Enabling Real-Time Analytics for Proactive Management

Functions like fraud detection, real-time credit scoring, and AML screening require real-time analysis of massive datasets. Traditional databases, optimized for OLTP, often lack the capability for efficient complex analytical queries on live data. This can lead to delays in identifying fraud or assessing credit risks, potentially resulting in substantial losses. A delay of seconds in fraud detection can cost millions.

AI in FinTech for fraud detection was valued at $12.2 billion in 2023 and is projected to reach $61.6 billion by 2032, with over 71% of financial institutions using AI to combat fraud.[1] Traditional batch processing for fraud detection is inadequate when funds can be illicitly transferred in moments. Today's approaches, often using Generative AI, synthesize varied datasets for risk mitigation, such as identifying early borrower distress or flagging implausible data inputs in real-time. This shift demands a database infrastructure that can serve both high-throughput transactional and complex analytical queries on the same fresh data without performance degradation.

Ensuring Data Consistency in Distributed Environments

FinTech companies must ensure strong data consistency across distributed systems, even with network partitions or hardware failures. Data consistency ensures all data remains accurate and uniform, preserving data integrity. Many distributed solutions compromise strong consistency for availability, which is unacceptable for financial systems where accuracy is paramount.

Inconsistencies can lead to faulty logic, erroneous calculations, failed audits, and eroded trust. Challenges include network latency, node failures, partition tolerance, and managing distributed transactions. Distributed system resilience is paramount for cloud-native FinTech applications. Incorrect balances or failed reconciliations can lead to severe penalties. The trend towards Open Banking, with payments projected to surge from $57 billion in 2023 to $330 billion by 2027 [1], amplifies the criticality of provable data consistency.

Achieving Elastic Scalability for Unpredictable Demand

FinTech platforms experience unpredictable fluctuations in traffic due to holiday sales, market volatility, or new product launches. Legacy systems relying on vertical scaling require manual intervention and downtime, leading to service disruptions and lost revenue.

FinTech applications must scale horizontally and vertically, adapting dynamically. Microservices and event-driven design are increasingly adopted. A payment FinTech overwhelmed by peak event traffic with on-premise servers migrated to a more scalable cloud infrastructure as an example. Elastic scalability is also about cost optimization, avoiding over-provisioning by scaling resources up and down rapidly in response to real-time demand.

Adhering to Strict Regulatory and Security Mandates

FinTech companies operate in a stringently regulated industry, facing scrutiny concerning data security, privacy, and compliance. 93% of FinTech professionals report struggling with compliance, and over 60% of companies incurred fines exceeding $250,000 last year.[1] The average cost of a data breach reached $4.88 million in 2024 [1], and 72% of finance Chief Risk Officers identify cybersecurity as their primary concern.

The financial sector was the most breached in 2023, with global cybercrime expenses projected to reach $10.5 trillion annually by 2025. Data privacy laws like GDPR impose significant obligations. 95% of business leaders view privacy as a business necessity, and 94% of customers would not engage with a company if their data was inadequately protected. A database solution must possess inherent features that simplify adherence to these mandates.

CHAPTER 2

Why TiDB is the Perfect Fit for High-Growth FinTech Applications

This chapter introduces TiDB, detailing its core architectural tenets and revealing why this distributed SQL database is uniquely positioned to address the rigorous demands of high-growth FinTech applications.

Horizontal Scalability

TiDB's ability to scale horizontally is critical for FinTech platforms. Its distributed architecture allows linear increases in transaction throughput and storage capacity by adding more nodes, eliminating bottlenecks associated with vertically scaled monolithic databases.

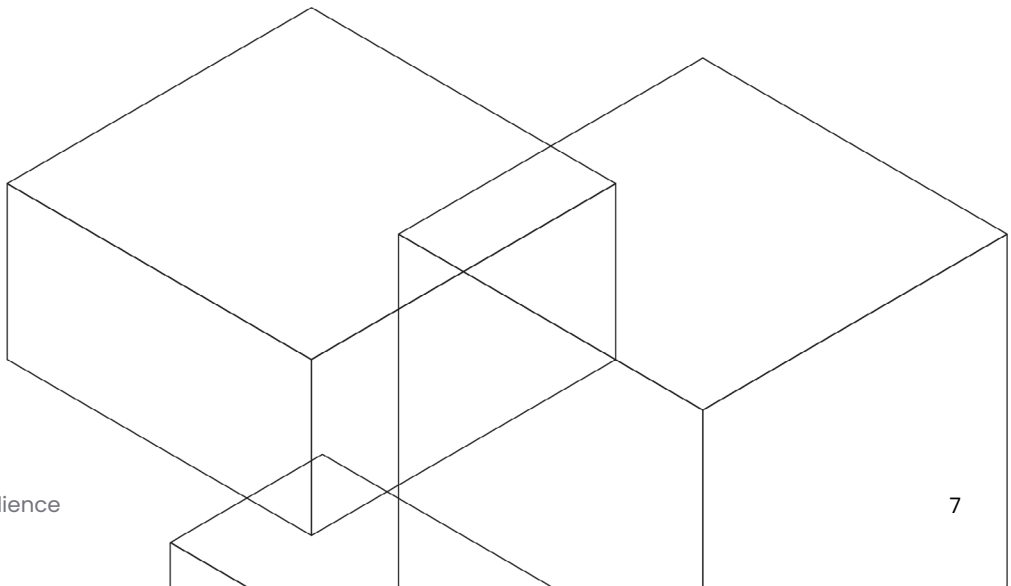

The TiDB platform comprises:

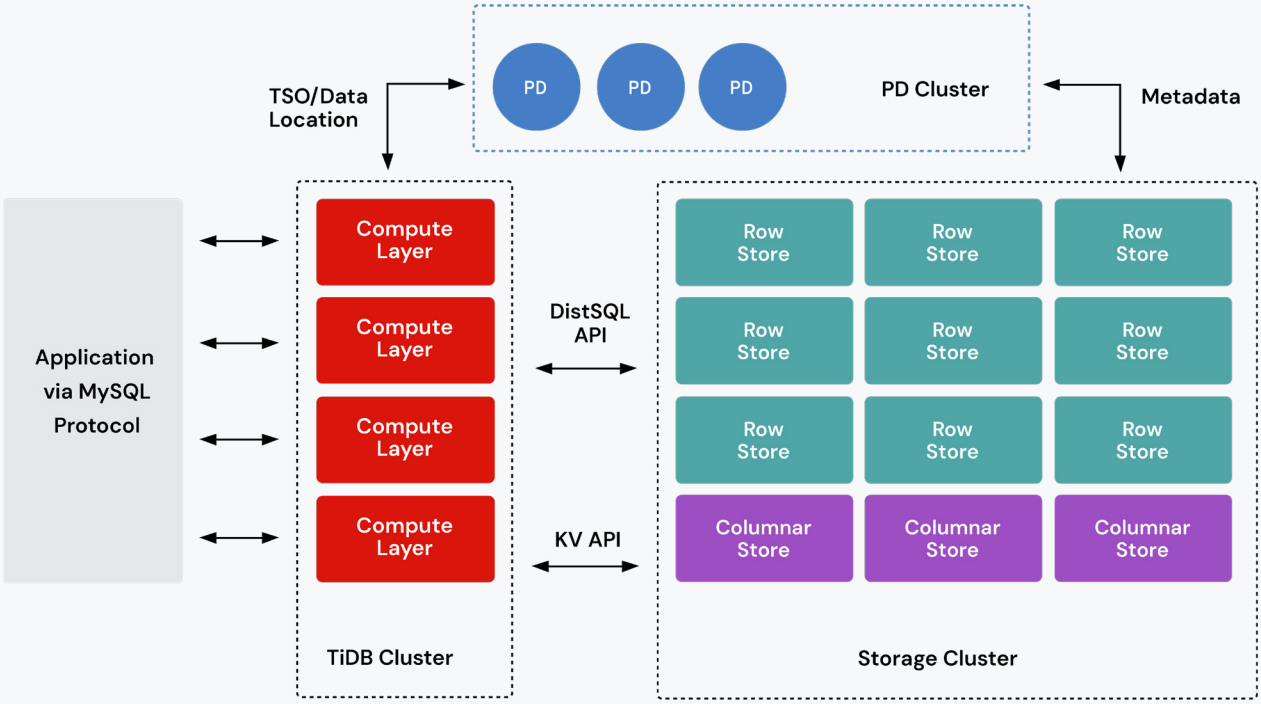

Figure 1. A diagram depicting TiDB's distributed SQL architecture that separates compute from storage.

- **Compute Layer:** A stateless SQL processing layer that parses queries, builds execution plans, and distributes tasks. Multiple TiDB servers can be deployed for scaling out query processing. Its MySQL compatibility eases migration.
- **Row-Based Storage:** A distributed transactional key-value storage engine. Data is sharded into "Regions," basic units of data replication using the Raft consensus protocol across storage nodes for redundancy and high availability. Regions can be split or moved as data grows, allowing horizontal storage scaling.
- **Metadata Manager:** Manages metadata, monitors storage nodes, balances Regions, and allocates globally unique timestamps for distributed transactions, crucial for consistency.

TiDB's decoupling of compute and storage layers allows independent scaling of compute and storage, enabling tailored resource allocation and optimizing performance and cost.

Mixed Workload Processing

TiDB's native support for mixed workload processing is powered by a columnar storage engine working alongside row-based storage. The columnar storage maintains asynchronous, real-time replicas of data from row storage using the Raft consensus protocol (as a Raft Learner) for data consistency. This architecture allows TiDB to serve OLTP and OLAP workloads from a single platform without ETL processes. Transactional data written to row-based storage is replicated to columnar storage in near real-time. Analytical queries can be directed to row-based storage for efficient processing due to its columnar format.

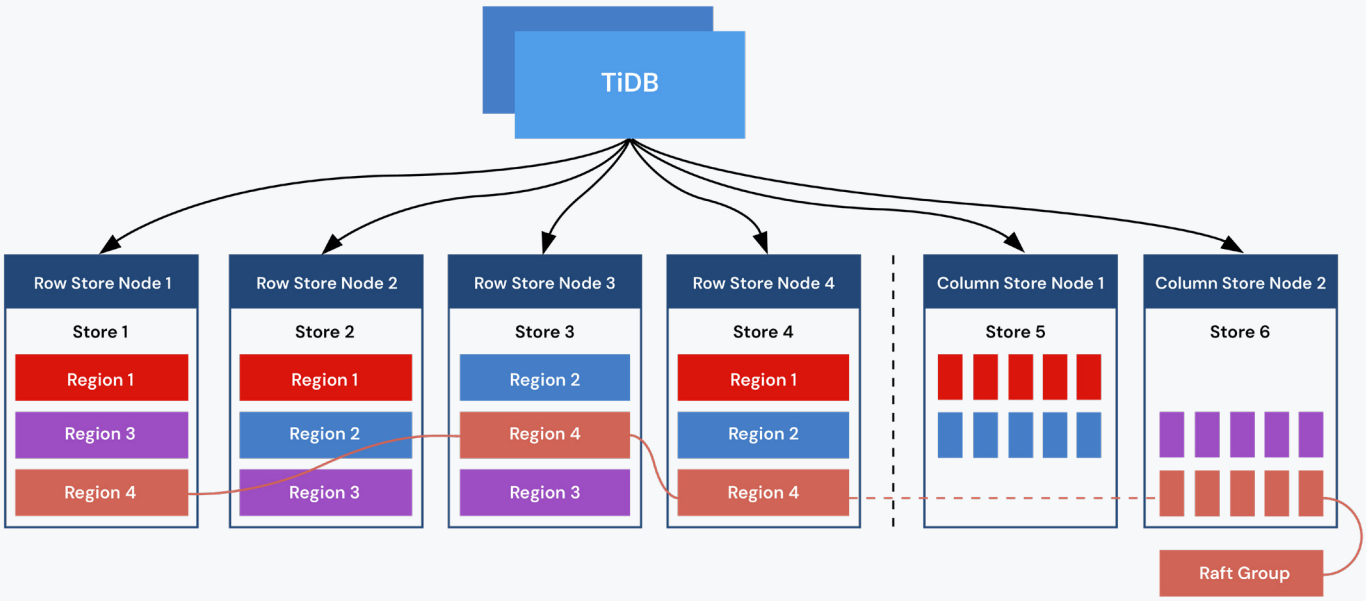

Figure 2. A diagram showing how TiDB handles mixed workload processing.

The TiDB optimizer intelligently routes query parts to row-based storage, columnar storage, or both. This enables real-time analytics on live transactional data for fraud detection, risk assessment, and personalized customer experiences, simplifying architecture and reducing operational overhead. This addresses the "data silos" problem, identified by 54% of financial institution leaders as a primary barrier to innovation.

Strong Consistency (ACID) and High Availability

Strong data consistency is non-negotiable for FinTech applications. As mentioned earlier, TiDB uses the Raft consensus protocol, ensuring data is reliably replicated and transactions adhere to ACID properties. TiDB provides Snapshot Isolation by default.

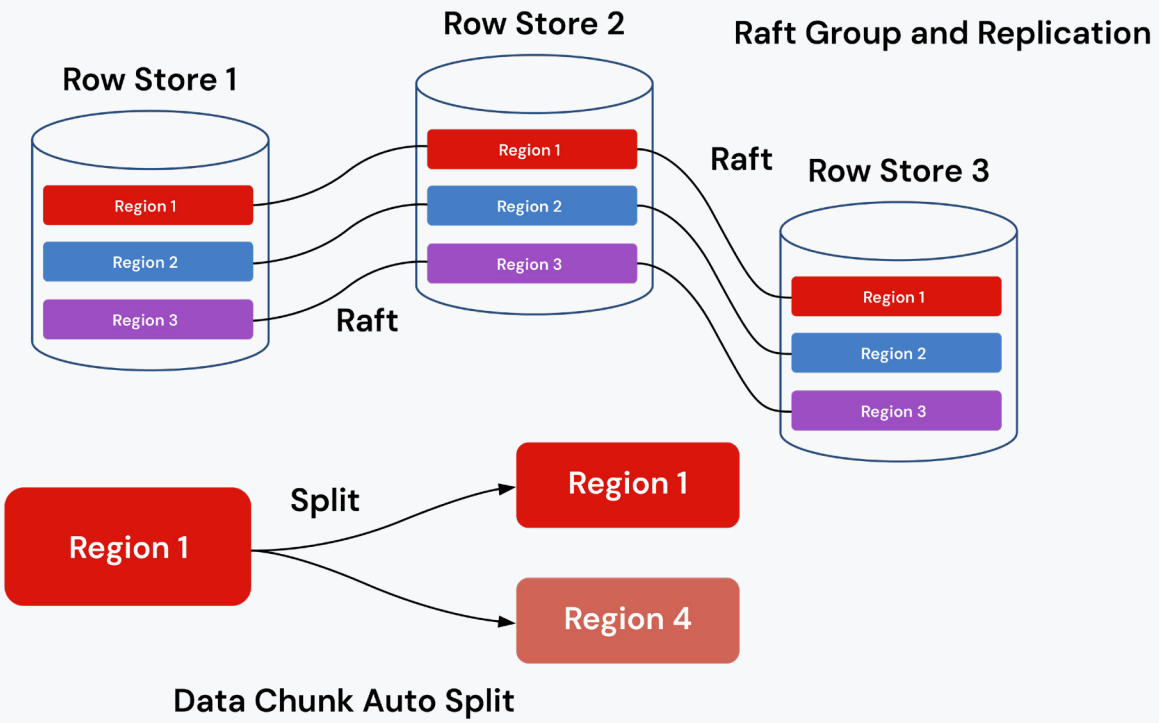

Figure 3. A diagram depicting how TiDB uses Raft for data replication.

Raft mandates writes are committed only after replication to a majority of replicas within a Region's Raft group, guaranteeing strong consistency even with network partitions or hardware failures. Automatic failover ensures high availability. TiDB clusters can be deployed across multiple data centers or cloud availability zones for robust disaster recovery. This combination of strong consistency and horizontal scalability is a key differentiator.

Cloud-Native Flexibility

TiDB is a cloud-native database, designed for dynamic infrastructure environments, offering deployment flexibility across public clouds, private clouds, on-premises data centers, or hybrid configurations. The TiDB Operator for Kubernetes automates lifecycle management of TiDB clusters on Kubernetes, including deployment, configuration, scaling, upgrades, backups, and failure recovery. This simplifies operations and allows adoption of DevOps practices. This cloud-native flexibility and Kubernetes automation reduce operational overhead, allowing teams to focus on core business logic and innovation.

CHAPTER 3

How TiDB Empowers FinTechs to Scale Faster and Smarter

Choosing the right database architecture is no longer just a technical decision. In FinTech, where data accuracy, uptime, and auditability are non-negotiable, your database must serve as both a performance engine and a control surface. It has to scale with your customer base, protect against systemic risk, and adapt to new products and regulatory shifts without disrupting day-to-day operations.

This chapter outlines the key architectural capabilities that distinguish a database built for FinTech. It also explains why distributed SQL—particularly TiDB—has emerged as a compelling foundation for long-term resilience and innovation.

Effortless Expansion: Elastic Scaling for Dynamic Workloads

TiDB's architecture, with decoupled compute, storage, and analytical layers, provides agility. This allows independent scaling of cluster aspects, enabling swift responses to demand shifts without disrupting live transactions or over-provisioning.

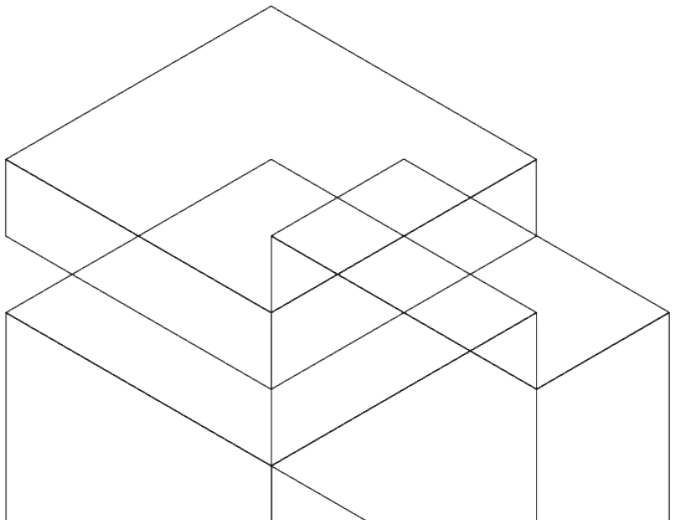

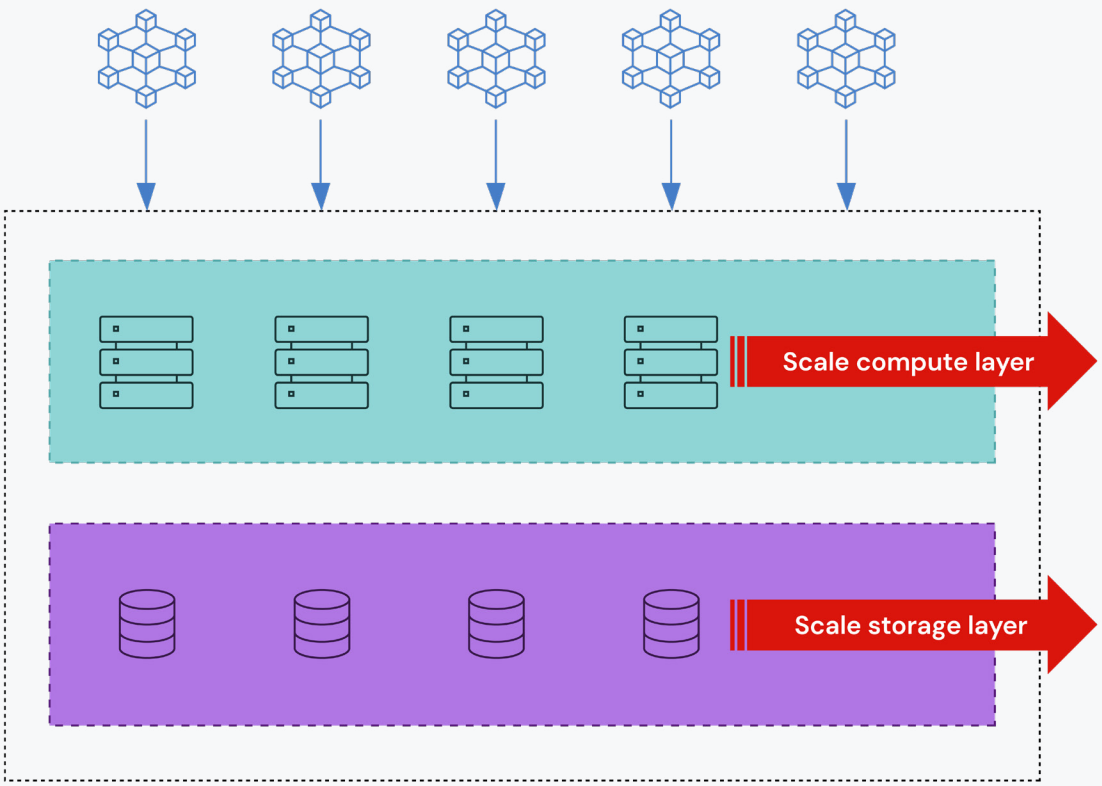

Figure 4. How TiDB's compute and storage nodes can be horizontally scaled independently without interruption.

For example, a peer-to-peer lending platform doubled its processing capacity in minutes during a funding surge by adding nodes. This "effortless expansion" significantly reduces time-to-market for scaling. Traditional database scaling can take weeks; TiDB scaling, especially with the TiDB Operator on Kubernetes, can take minutes or hours. This elastic scaling democratizes handling "viral" growth, allowing startups to scale massively and cost-effectively.

Accelerated Decision-Making: Leveraging Mixed Workload Processing for Real-Time Advantage

TiDB's mixed workload processing capabilities enable faster, informed decisions by facilitating real-time analytics on live transactional data for fraud prevention, credit underwriting, personalized offers, and risk assessment. A digital bank reduced fraud response time from hours to milliseconds by running transactional queries and ML fraud detection models on the same TiDB cluster. This shift from reactive to proactive intervention curtails financial losses and strengthens security. Insights from mixed workload processing can fuel product improvement and enhance customer experience by analyzing real-time behavior, leading to more tailored services and fostering deeper customer engagement.

Optimized Compliance: Achieving Cost-Effective Regulatory Adherence

TiDB alleviates the burden of navigating complex regulatory landscapes through its architectural properties and built-in features. Strong ACID compliance ensures data integrity and auditable transaction trails. Features like robust backup/recovery (including PiTR) and data encryption simplify adherence to data protection mandates.

Consolidating multiple legacy systems into a single TiDB cluster can yield significant cost savings. A payment processing company reported a 40% reduction in compliance-related costs after migrating to TiDB. This consolidation reduces complexity in managing data and generating reports. 93% of FinTechs report compliance struggles.[1] TiDB streamlines these processes, allowing resource redirection towards proactive risk management and innovation.

How TiDB Solves Specific FinTech Challenges

	TiDB's Advantage	Key Enabling Technologies/Features
Handling High-Volume Transactions	Horizontal Scalability & High Throughput	Distributed Architecture, Raft, Load Balancing
Real-Time Analytics for Risk Management	Mixed Workload Processing	Columnar Storage, Real-time Replication, SQL Compatibility
Data Consistency Across Distributed Systems	Strong Consistency (ACID Compliance) & High Availability	Raft, Distributed Transactions
Elastic Scalability for Unpredictable Workloads	Decoupled Compute/Storage & Cloud-Native Elasticity	Independent Scaling, Kubernetes Operator
Adhering to Regulatory & Security Mandates	Unified Architecture, Built-in Security, Auditable & Consistent Data	Encryption, Backups, Point-in-Time Recovery, Strong Consistency

CHAPTER 4

Real-World Success Stories

From global FinTech disruptors to payment processors and tier-one banks, leading financial services organizations are rearchitecting their data platforms to meet rising demands for scalability, real-time responsiveness, and developer efficiency. The following companies chose TiDB to eliminate architectural bottlenecks, future-proof performance, and simplify operations without sacrificing consistency or control.

Reduces Operational Overhead 25% and Eliminates Unexpected Downtime

Plaid, a global leader in financial data connectivity, helps thousands of apps and institutions securely access consumer financial data. As usage grew, Plaid's MySQL-based architecture running on Amazon Aurora introduced scaling limitations and developer friction.

The Challenge

- Manually sharded clusters created high operational overhead
- Schema changes were brittle and time-consuming
- Cross-tenant observability and query routing were complex
- Development velocity slowed under infrastructure constraints

The Solution

Plaid adopted TiDB to eliminate manual sharding and accelerate infrastructure scalability. With TiDB's MySQL compatibility, strong consistency via Raft, and native horizontal scalability, Plaid migrated a greenfield service to TiDB on AWS using Kubernetes. This allowed seamless scaling, faster development cycles, and unified observability—without compromising performance.

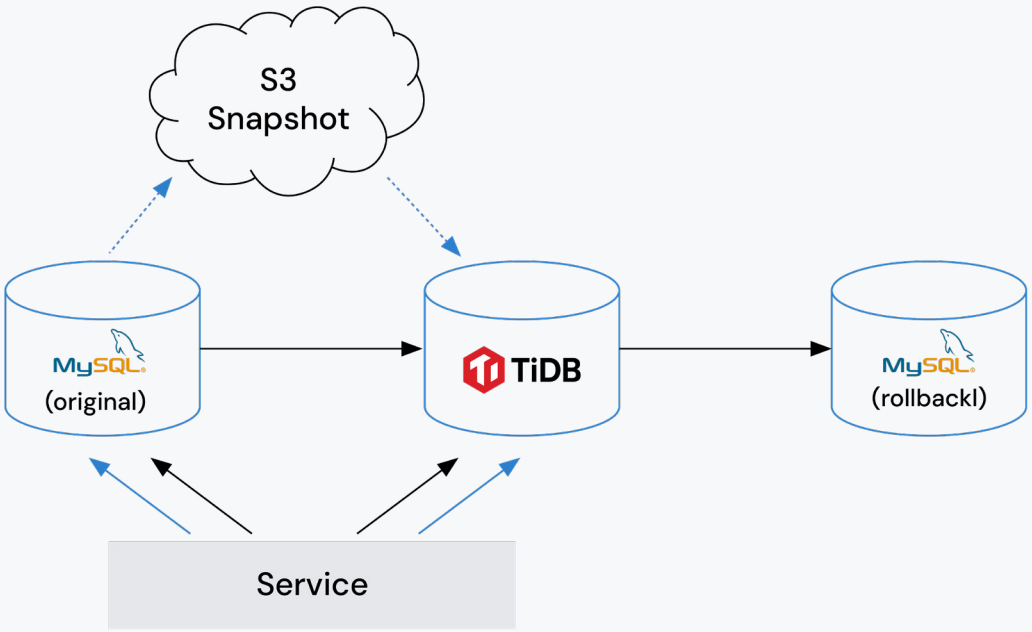

Figure 5. Plaid's approach to migrating a new service to TiDB instead of retrofitting an existing one.

THE RESULTS

- **25% reduction in operational overhead** through elimination of manual sharding
- **5x increase in resource utilization** thanks to horizontal elasticity and simplified topology
- **Zero-downtime schema changes** using automation and TiDB's resilient design
- **Faster engineering velocity** due to centralized queries, simplified backups, and streamlined development workflows
- **Optionality for future real-time analytics** without needing a separate OLAP system

> *"With TiDB, we can now perform upgrades with zero downtime alongside large table schema migrations." —* PLAID

Watch the customer testimonial video »

Square

Building a Real-Time, Multi-Tenant Financial Platform

As a global financial technology leader, Square processes billions in payments across millions of merchants and consumers. With a growing portfolio of embedded financial products, Square required a resilient data infrastructure that could support real-time processing, tenant isolation, and regulatory-grade consistency across multiple regions.

The Challenge

- Scaling transactional workloads across diverse services and markets
- Eliminating the complexity of sharding and schema duplication
- Ensuring high availability and data consistency in a distributed setting
- Supporting both transactional and analytical queries with minimal overhead

The Solution

Square adopted TiDB to consolidate previously fragmented MySQL workloads into a horizontally scalable, strongly consistent architecture. TiDB's compatibility with MySQL allowed Square to migrate services with minimal code changes while achieving better performance, simplified operations, and cross-tenant observability.

THE RESULTS

- **Improved developer velocity** through centralized schema management and online DDL
- **Zero-downtime scaling** to meet seasonal and regional demand spikes
- **Operational simplicity** with unified data access and real-time metrics
- **Resilience and reliability** in a multi-region deployment setup

TiDB helped Square deliver consistent, high-performance experiences to customers and merchants alike—while simplifying infrastructure at scale.

> *"With TiDB, we no longer have to worry about the size of our database. TiDB just runs quietly and we don't worry about it." —* Square

Watch the entire tech talk »

SB Payment Service

Achieving Zero-Downtime Scaling

SB Payment Service, a major payment processor in Japan, supports hundreds of secure financial transactions per second. As digital payments surged, their MySQL-based backend showed limitations in high availability, peak performance, and failover handling.

The Challenge

- Ensuring seamless performance under high concurrency
- Managing failover and updates without disrupting service
- Reducing complexity while scaling infrastructure securely

The Solution

SB Payment Service adopted TiDB for its new online payment system. Key factors were TiDB's distributed SQL capabilities, horizontal scalability, strong consistency, and high MySQL compatibility, which minimized migration effort. A rigorous Proof of Concept (PoC) included performance and failure tests.

THE RESULTS

- **Zero-Downtime Scaling and High Availability:** Achieved continuous service even during peak periods and updates
- **Enhanced Performance:** Demonstrated processing performance approximately 1.7 times greater than their previous MySQL system under equivalent hardware
- **Reduced Recovery Time:** Significantly reduced downtime for recovery during failures or maintenance
- **Operational Efficiency:** TiDB's management console was easy to use, with most tasks achievable with a single click
- **Strategic Platform:** TiDB is now a core database technology for SB Payment Service's future development. PingCAP's support was highly satisfactory

> *"The ability to drastically reduce the read replica switch time is a clear strength of TiDB as a distributed SQL database." —* SB Payment Service

Read the full case study »

These success stories highlight how leading FinTech organizations are overcoming legacy infrastructure limitations by embracing distributed SQL. Whether it's eliminating sharding complexity, ensuring zero-downtime scalability, or powering real-time transaction processing at global scale, TiDB enables financial platforms to operate with greater resilience, agility, and confidence. As data demands continue to rise, these companies demonstrate what's possible when the database becomes a driver—not a bottleneck—of innovation.

CHAPTER 5

Best Practices for Deploying TiDB in FinTech Applications

To maximize TiDB's benefits, FinTechs should adhere to proven best practices for deployment, performance tuning, and monitoring.

Deployment Strategies: Cloud-Native, On-Premises, and Hybrid Approaches

Choosing the right deployment strategy is critical and should align with business model, technical capabilities, risk appetite, and regulatory obligations.

- **Cloud-Native Deployment:** For cloud-first strategies, deploying TiDB on Kubernetes using the TiDB Operator is recommended for automation of lifecycle management tasks like deployment, scaling, upgrades, monitoring, and backup/restore. This aligns with DevOps practices.

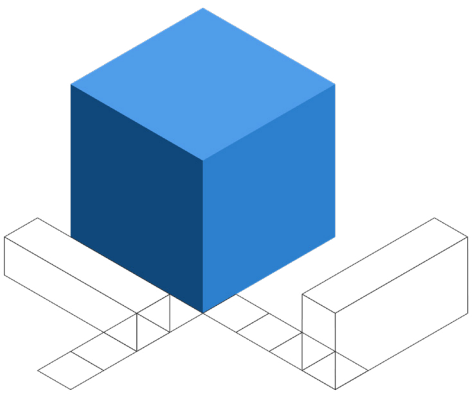

- **On-Premises/Hybrid Deployment:** For significant on-premises infrastructure or specific data sovereignty needs, TiDB can be deployed in private data centers or hybrid configurations. Careful planning of network topology (low latency, sufficient bandwidth) and use of enterprise-grade SSDs are crucial. For hybrid deployments, robust network connectivity and latency management are essential.

Tip: Choose a deployment model—cloud, hybrid, or on-prem—that aligns with your regulatory, data residency, and risk management requirements.

Performance Tuning Tips

Achieving peak performance involves appropriate hardware, schema design, and query optimization.

- **Strategic Indexing:** Implement composite indexes for frequent queries. Analyze execution plans regularly. Define primary keys based on common access patterns.
- **Effective Partitioning:** For large tables, use TiDB's table partitioning (e.g., by date or region) to improve query performance, optimize I/O, and simplify data management.
- **SQL Query Optimization:** Craft efficient SQL queries, minimizing rows scanned and selecting optimal join strategies. Be mindful of write hotspots.
- **Hardware Considerations:** Use enterprise-grade SSDs for row store nodes. For column store nodes, ensure CPUs support required instruction set extensions (AVX2 for x86-64, ARMv8 for ARM64) for optimal analytical performance. Performance tuning is an ongoing process.

Tip: Design indexes based on real query patterns and continuously refine them using TiDB's built-in performance insights..

Observability and Monitoring Frameworks

Comprehensive observability and monitoring are critical for health, performance, and reliability.

- **Metrics Collection and Visualization:** Integrate TiDB with Prometheus and Grafana. TiDB exposes metrics on query latency, QPS, node health, Raft status, storage utilization, and transaction statistics.
- **TiDB Dashboard Utilization:** Leverage the built-in TiDB Dashboard for analyzing slow queries, identifying bottlenecks, viewing cluster topology, and capacity planning.
- **Network Monitoring:** Monitor network latency and bandwidth between TiDB components, especially between row store nodes and across data centers.
- **Alerting Mechanisms:** Implement alerting for critical metrics like high latency, low disk space, or node failures.
- **Data Consistency Checks:** Regularly perform data consistency checks using TiDB's built-in commands (e.g., ADMIN CHECK TABLE). Observability is fundamental for capacity planning, cost management, operational risk management, and regulatory compliance.

Tip: Enable observability from day one to support SLA monitoring, performance tuning, and compliance readiness.

Successfully deploying TiDB in FinTech environments requires more than a fast database. It demands architectural discipline, regulatory awareness, and operational foresight. By following proven best practices in deployment design, indexing strategy, and observability, engineering teams can unlock TiDB's full potential while maintaining the resilience, precision, and compliance financial platforms require. As workloads scale and expectations rise, these foundations ensure that your data infrastructure stays ready for what's next.

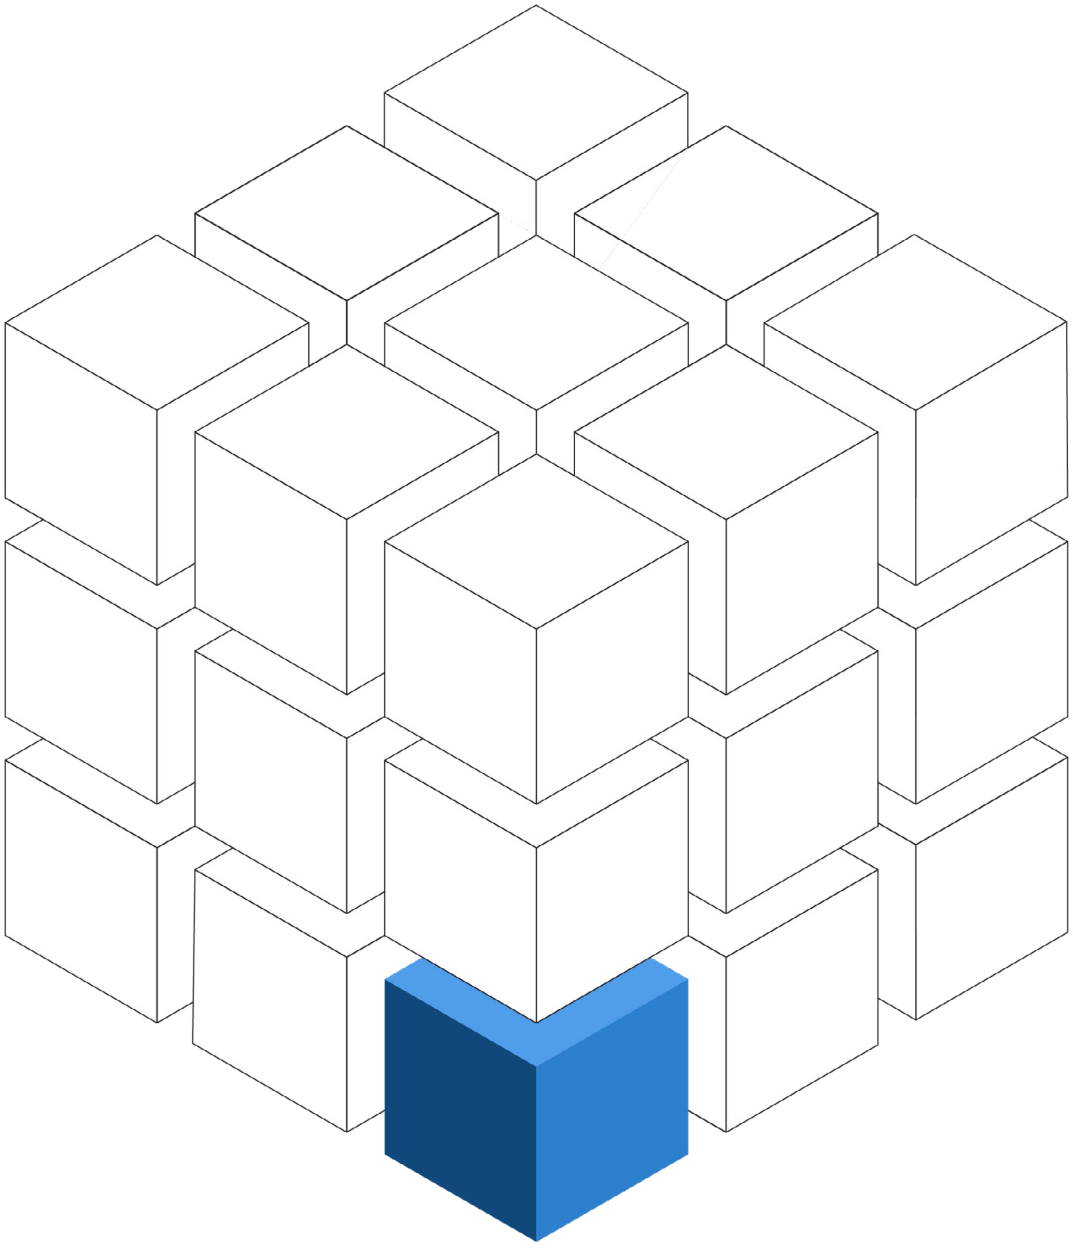

Conclusion

TiDB is a strategic data foundation engineered to empower FinTech organizations to navigate the complexities of finance. By addressing challenges of scale, real-time processing, data consistency, and regulatory compliance, TiDB enables FinTechs to innovate faster, operate efficiently, deliver superior customer experiences, and architect the next generation of financial services.

TiDB's unique architecture—featuring horizontal scalability, mixed workload processing, strong ACID consistency through Raft, and cloud-native flexibility with Kubernetes—provides a comprehensive solution to traditional database limitations. This translates to accelerated innovation, reduced operational costs, and future-proofed data infrastructure.

As the FinTech industry evolves, data infrastructure demands will intensify. TiDB is architected for these future challenges, positioning it as a critical enabler of growth and innovation. Adopting TiDB can catalyze broader digital transformation, unlocking new business models and service offerings. The open-source nature of TiDB fosters continuous improvement and collaborative innovation, offering FinTechs greater control and flexibility. Partnering with TiDB means investing in a future-ready data foundation to support building the future of finance.

Ready to scale your financial platform to the next level? Schedule a personalized tech talk with one of our database experts to explore how distributed SQL can help you scale securely, meet compliance needs, and accelerate innovation within your company.

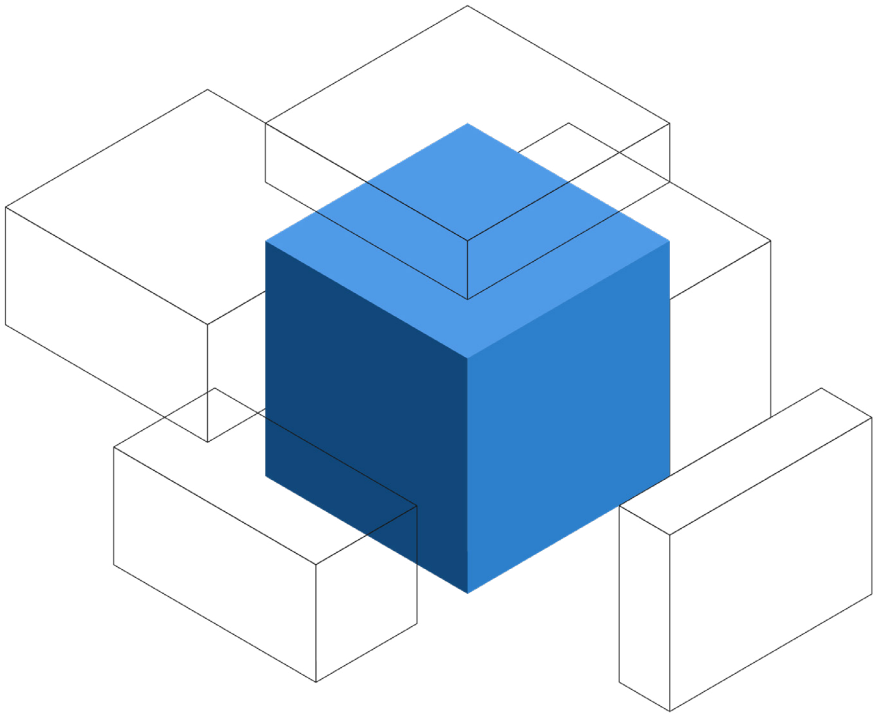

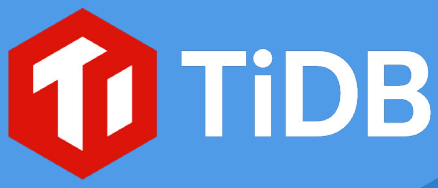

EVALUATE TiDB FOR YOURSELF

Start Your Free Trial

Contact us for a personalized demo at pingcap.com/demo/