



Your Path to *Scalable* Fintech Resilience



PLAID



Square



SB Payment Service

ele-style

motion



banking

Use Cases

- **Payments that stay online:**

Checkout and payout windows miss SLAs during bursts. TiDB's online DDL and rolling upgrades keep ledgers live while scaling reads and writes horizontally.

- **Ledgers you can evolve safely:**

Schema changes block releases or require weekend freeze windows. TiDB adds columns/indexes online and rolls forward without app rewrites.

- **Multi-tenant control without silos:**

One-DB-per-tenant explodes cost; noisy neighbors break SLOs. TiDB's Resource Groups cap per-tenant RU/IO, and Placement Rules pin hot data to fast nodes and keep cold data cheap

- **Risk & analytics on fresh data:**

ETL lag forces stale fraud signals and duplicate stores. TiDB's columnar data store runs analytics on live OLTP data while streaming events to models and data platforms.

Trusted and Verified by Innovation Leaders

“ With TiDB, we can now perform upgrades with zero downtime alongside large table schema migrations.

– Zander Hill, Platform Architecture Lead, Plaid

Discover the TiDB Advantage:


- **Open source by design:** Transparent roadmap, community-driven innovation, and no proprietary lock-ins
- **Zero-downtime upgrades:** Rolling upgrades plus online schema evolution mean you change prod while traffic flows
- **Elastic writes, not just reads:** Scale both reads and writes, so you can retire manual sharding
- **MySQL compatibility:** Keep drivers/ORMs and operational muscle memory as you grow

Contact our experts today to schedule an intro meeting. Email us at contact@pingcap.com.

How Does TiDB Work?

TiDB's distributed SQL architecture is made up of three components:

- 1 SQL layer:** A stateless SQL layer that's MySQL compatible and decouples compute from storage to simplify scaling. Each node can take both reads and writes.
- 2 PD cluster:** A cluster manager that dynamically balances the data load in real time while mitigating hotspots and providing for the implementation of customized scheduling policies.
- 3 Storage layer:** A distributed storage engine that offers built-in high availability and strong ACID consistency that can auto-scale to hundreds of nodes and petabytes of data.



TiDB Success Story



This leading fintech company builds a data transfer network for applications to connect with users' bank accounts. The company incurred unexpected system downtime in 2022 due to Amazon Aurora and a later forced upgrade from MySQL 5.7 to 8.0.

After exploring several potential new solutions, Plaid chose TiDB for its reliability, lower costs resulting in 20% less FTE resources than Amazon Aurora, improved velocity for developers to make online schema changes and upgrades in production, and lower TCO as 270 Amazon Aurora clusters were consolidated into 5 TiDB clusters.

Amazon Aurora's architecture forced Plaid to overprovision resources, leading to low utilization. For the same cost as Amazon Aurora, TiDB now allows the company to avoid the cost of future planned and unplanned downtime, acquire new customers, and future-proof their SQL as the business continues to grow. For more details, check out the customer story: pingcap.com/customers/



Follow Us

