



Unlocks Limitless Scale for Data-Intensive Businesses

What is TiDB?

With over 39,000 GitHub stars and a growing community of contributors, TiDB is an advanced open-source, distributed SQL database. At its core, TiDB combines the scale of NoSQL databases with the full functionality of a single-node relational SQL database.

- **Zero-downtime by design:** Raft-based HA with ACID transactions, online DDL, rolling upgrades, and PITR.
- **MySQL-compatible, AI-ready:** Drop in with existing MySQL tooling and skills, then add AI-centric capabilities (branching, vector/semantic queries) without a rewrite.
- **Adaptive scale:** Decoupled compute + storage with automatic sharding and placement lets you create databases in seconds and scale to thousands of instances as workloads spike.
- **Unified storage engine:** Row and columnar storage under one planner for fresh OLTP + analytics, with optional vector similarity/semantic querying in the same platform.
- **Built for multi-tenancy & agents:** Per-tenant isolation, resource groups, and copy-on-write branching enable safe parallel experiments and schema evolution without impacting neighbors.

TiDB Use Cases

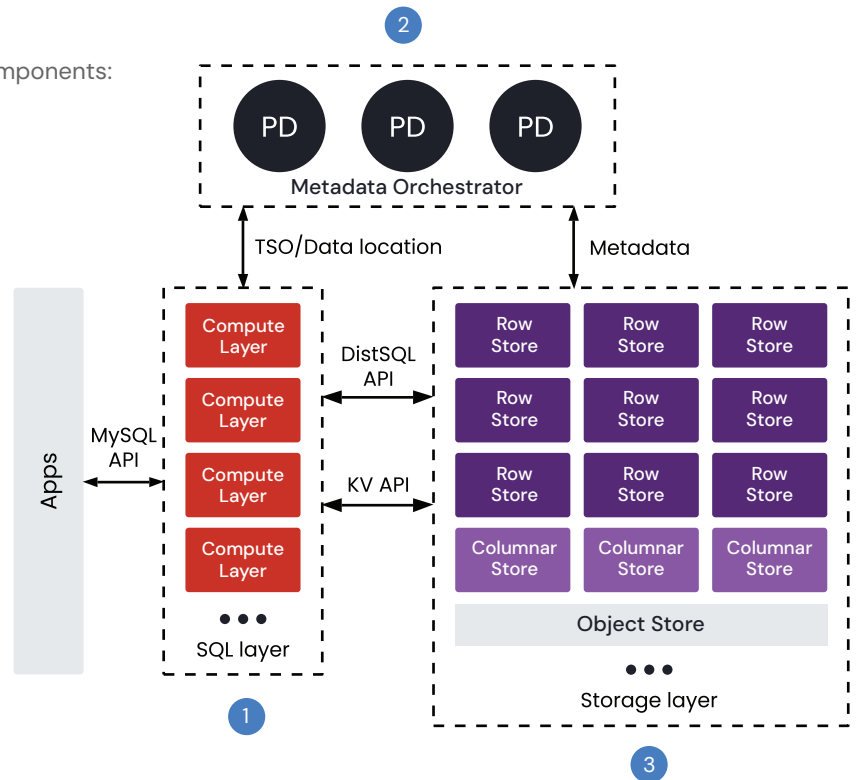
TiDB excels in scenarios that require strong ACID consistency, 24/7 availability, extreme scalability of TBs - PBs, high read/write throughput, or thousands to millions of QPS:

- **A more scalable and reliable RDBMS:** Drop-in MySQL compatibility with horizontal scale, automated failover, PITR, and zero-downtime operations.
- **Database consolidation:** Replace fleets of siloed RDBMS instances for higher dev velocity, lower ops toil, and better resource efficiency.
- **Operational data layer:** Process and query large volumes of fresh data with high throughput and low latency.
- **NoSQL replacement:** Keep NoSQL-like elasticity while retaining SQL, ACID transactions, secondary indexes, and rich joins—no app rewrites.
- **Agentic AI apps:** Spin up thousands of isolated databases in seconds for agent-built apps.
- **Real-time RAG & vector search:** Serve OLTP on fresh data while running analytics and vector/semantic queries in the same platform, eliminating fragile ETL hops.

How Does TiDB Work?

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

- 1 Compute layer:** A stateless SQL layer that decouples compute from storage to simplify scaling. Each node can take both reads and writes.
- 2 Metadata orchestrator:** 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 backed by shared object storage that delivers built-in high availability and strong ACID consistency while auto-scaling to hundreds of nodes and petabytes of data.



TiDB Success Stories

For additional insights, please visit pingcap.com/customers.

- Plaid:** Migrated key services from Amazon Aurora to TiDB to retire manual sharding, run online schema changes with zero downtime, and scale elastically on Kubernetes. Result: simpler ops, faster iteration, and headroom to grow without re-architecting.
- Databricks:** TiDB saved Databricks over 60% of their TCO because of TiDB's horizontal scalability, decoupled compute and storage, and automated database operations. With TiDB, Databricks never experienced downtime across all three public clouds—AWS, GCP, and Azure—in over two years.
- Pinterest:** TiDB helped Pinterest cut their TCO by 80%–50% reduction in hardware costs, 30% reduction in engineering costs—while eliminating complex systems. With TiDB, Pinterest can now support hundreds of microservices querying three billion unique products at ultra low latency, delivering a much improved user experience.
- Atlassian:** TiDB consolidated thousands of fragmented Postgres/DynamoDB instances into a single distributed SQL backbone across 13 regions, cutting instance sprawl to ~1% and lowering vCPU. With online DDL and rolling operations, Atlassian sustains 99.99% uptime at petabyte scale and 3M+ tables.



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. With AI-driven innovations and multi-cloud flexibility, TiDB offers unmatched agility, resilience, and security.

440 N Wolfe Road, Office EL265
Sunnyvale, CA 94085

www.tidb.io
contact@pingcap.com

SEQUOIA | GGV CAPITAL

COATUE | ACCESS
TECHNOLOGY VENTURES

LinkedIn: <https://www.linkedin.com/company/pingcap/>
Twitter: <https://twitter.com/PingCAP>
YouTube: <https://www.youtube.com/c/PingCAP>