

YugabyteDB is an open-source, high-performance distributed SQL database for building global, cloud-native applications. It can run anywhere cloud-native applications are deployed, across private, public, hybrid, and multi-cloud environments, on VMs, containers, or bare metal. The database combines powerful RDBMS capabilities with cloud-native resilience, scale, and geo-distribution. Most importantly, YugabyteDB eliminates forced trade-offs between different databases by delivering capabilities organizations need in their cloud-native databases:

- Widely adopted SQL API
- RDBMS capabilities, including distributed ACID transactions, stored procedures, triggers, and more
- Resilience and continuous availability
- Horizontal scalability
- Geo-distribution
- Deployment in any cloud

YugabyteDB enables enterprises to build resilient systems of record that scale horizontally without compromising performance or integrity. The database is feature, wire, and code compatible with PostgreSQL. YugabyteDB reuses PostgreSQL's query layer to achieve a high degree of compatibility with existing PostgreSQL applications or those that can be migrated to PostgreSQL. With YugabyteDB, critical services can remain available during node, zone, region, and data center failures with fast failovers. Operations teams can effortlessly scale out even under heavy load without disruption or downtime by simply adding nodes to their YugabyteDB cluster. YugabyteDB offers the most comprehensive and flexible deployment options in geo-distributed environments, including synchronous and asynchronous replication as well as geo-partitioning.

YugabyteDB is developed and distributed as an [Apache 2.0 open source project](#).

YugabyteDB brings together four must-have needs of cloud-native apps, namely: SQL as a flexible query language, low-latency performance, continuous availability, and globally distributed scalability. Other databases do not serve all four of these needs simultaneously.

- Monolithic SQL databases offer SQL and low-latency reads but neither have the ability to tolerate failures nor can scale writes across multiple nodes, zones, regions, and clouds.
- Distributed NoSQL databases offer read performance, high availability, and write scalability but give up on SQL features such as relational data modeling and ACID transactions.
- Cloud-provider databases impose limitations, lock-in, and can be expensive at scale.

Additional benefits of **YugabyteDB** include:

Accelerate speed to market

YugabyteDB helps organizations accelerate time to market for their apps by standardizing on a single, distributed database to support a multitude of workloads requiring both SQL and NoSQL capabilities.

Reduce revenue impact due to downtime

YugabyteDB is designed to survive multiple failures from nodes to zones and even cloud regions and data centers.

Globalize data service deployment with strong data consistency

YugabyteDB allows distributing data between regions and clouds with strong acid consistency.

Horizontal scalability to deliver efficiency

YugabyteDB offers the ability to horizontally scale nodes globally, across regions, while remaining fully ACID compliant. YugabyteDB also delivers linear write/read scalability, plus auto-rebalancing, partitioning, and failover.

Be cloud-native and agnostic for future-proofing

Administrators and operators can scale YugabyteDB on-demand, across any infrastructure without downtime with a push-button deployment/scale model with no additional operational complexity.

Innovate with familiar APIs to mitigate the fear of change

YugabyteDB's API compatibility is aimed at accelerating developer onboarding. By integrating well with the existing ecosystem, YugabyteDB ensures that developers can easily get started using a language they are already comfortable with.