yugabyteDB

6 GEO-DISTRIBUTED DEPLOYMENT OPTIONS YOU WANT TO CONSIDER

Businesses build and deploy globally-distributed applications to meet customer expectations for always-on and highly responsive access. A geo-distributed database is a key component of this strategy, providing low-latency access to data and making data service resilient to cloud failures.

A YugabyteDB cluster consists of 3 or more nodes. They communicate with each other and distribute data among themselves.

MULTI-ZONE CLUSTER

Deploy the nodes of a YugabyteDB cluster in different zones within the same region.

Benefits:

- Resilient to a zone failure
- High availability
- Strong consistency
- Low read and write latency within the region

Tradeoffs:

- Higher read/write latencies for remote regions
- Lack of resilience to region-level outages, such as natural disasters

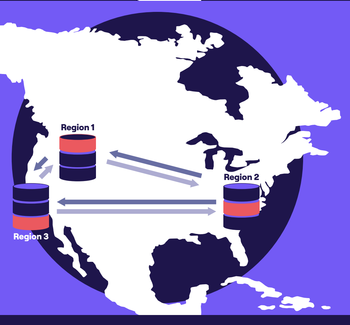

MULTI-REGION "STRETCHED" CLUSTER

Deploy the nodes of a YugabyteDB cluster in different regions.

Benefits:

- Resilient to a region failure
- High availability
- Consistent writes; tunable reads
- Low read and write latency within the region

Tradeoffs:

- Write latency can be high depending on distance and/or network packet transfer times
- Follower reads trade consistency for latency

MULTI-REGION CLUSTERS WITH UNIDIRECTIONAL ASYNCHRONOUS REPLICATION

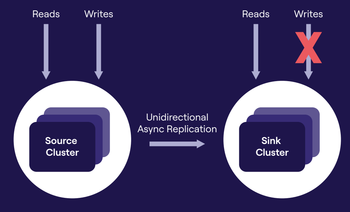

Deploy YugabyteDB cluster across two data centers or regions with asynchronous replication (active-passive configuration)

Benefits:

- Resilient to a zone failure, if nodes of each cluster are deployed across zones
- Strong consistency in source cluster; timeline consistent in sink cluster
- Low read and write latency within the source cluster region

Tradeoffs:

- High latency for clients outside source cluster region
- Database triggers won't get fired since xCluster bypasses the replicated records query layer

MULTI-REGION CLUSTERS WITH BI-DIRECTIONAL ASYNCHRONOUS REPLICATION

Deploy YugabyteDB cluster across two data centers or regions with asynchronous replication (active-active configuration)

Benefits:

- Resilient to a zone failure, if nodes of each cluster are deployed across zones
- Strong consistency in source cluster; timeline consistent in sink cluster
- Low read and write latency within either cluster

Tradeoffs:

- Database triggers won't fire since xCluster bypasses query layer, potentially leading to unexpected behavior
- Conflicting writes in separate universes can violate unique constraints and cause inconsistencies in the table and index
- The active-active mode does not support auto-increment IDs, so UUIDs are recommended instead

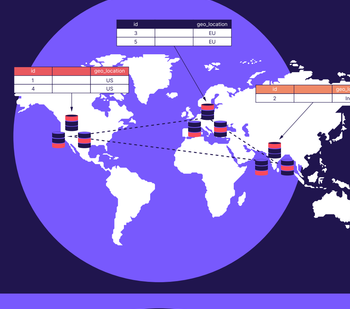

ROW-LEVEL GEO-PARTITIONING WITH DATA PINNING

Pin data to regions for compliance and lower latencies

Benefits:

- Resilient to a zone failure, if nodes of each cluster are deployed across zones
- Strong consistency
- Low latency within region; high latency across regions

Tradeoffs:

- Best suited for datasets that can be logically partitioned
- Accessing pinned data from outside the region incurs cross-region latency

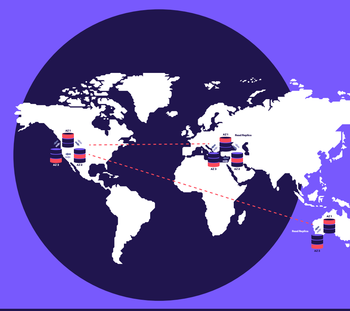

READ REPLICAS

Replicate data asynchronously to one (or more) read replica clusters

Benefits:

- Resilient to a zone failure, if nodes of primary cluster are deployed across zones
- Strong consistency in primary cluster; timeline consistent in replica clusters
- Low latency within region; write latency across regions dependent on distance

Tradeoffs:

- The primary cluster and read replicas are correlated, so adding read replicas does not improve resilience
- Read replicas cannot accept writes, leading to high write latency from remote regions even with a nearby read replica

yugabyteDB

yugabyte.com | contact@yugabyte.com

Learn more:

Geo-Distribution in YugabyteDB: Engineering Around the Physics of Latency