


# Cockroach DB: Distributed SQL for a Cloud Native Future

Distributed SQL delivers resilient, horizontal scale with always-on availability and data partitioned by location.



## CLOUD NATIVE

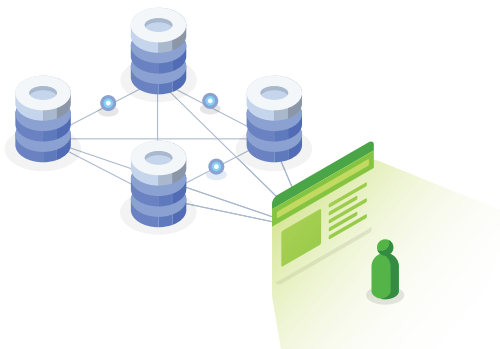
### Purpose-Built for the Cloud & Kubernetes

- Delivers on the key cloud-native primitives
- Enables stateful workloads on Kubernetes
- Deployable in any environment

## CLOUD SCALE

### Scale Your Data, Not Operational Complexity

- Scale by simply adding new nodes to a CockroachDB cluster
- Automate balancing/distribution of data to eliminate manual sharding
- Optimize server utilization evenly across all nodes



## ULTRA-RESILIENT

### Access Your Data, All the Time & Everywhere

- Eliminate complex, expensive backup configurations
- Avoid arduous, annoying remediation events (Zero RPO)
- Rolling upgrades eliminate downtime & ease updates/schema changes

## HYBRID-CLOUD

## Architect Your Data However You Want

- Eliminate the need to fork lift all data to the cloud at once
- Deploy on-prem & in the cloud and partition data across both
- Decompose your monolith at your own pace



## HIGHER ORDER OF CONSISTENCY

## Serializable Isolation in Distributed SQL

- Enjoy full ACID transactions at scale in a distributed environment
- Guarantee serializable isolation in cloud-neutral database
- No discrepancies or 'eventual' resolutions

## GEO-PARTITIONING

## Give Your Data a Location

- Ensure low latency to end-users wherever they are
- Pin data to a locale to comply with data protection laws
- Tie specific data to specific clouds or datacenters



# CockroachDB is Your Bridge to the Cloud

CockroachDB is helping companies **modernize legacy systems** and launch **net new cloud-native applications**.

