
RoachFest24

CockroachDB enables resilient, fully consistent clusters

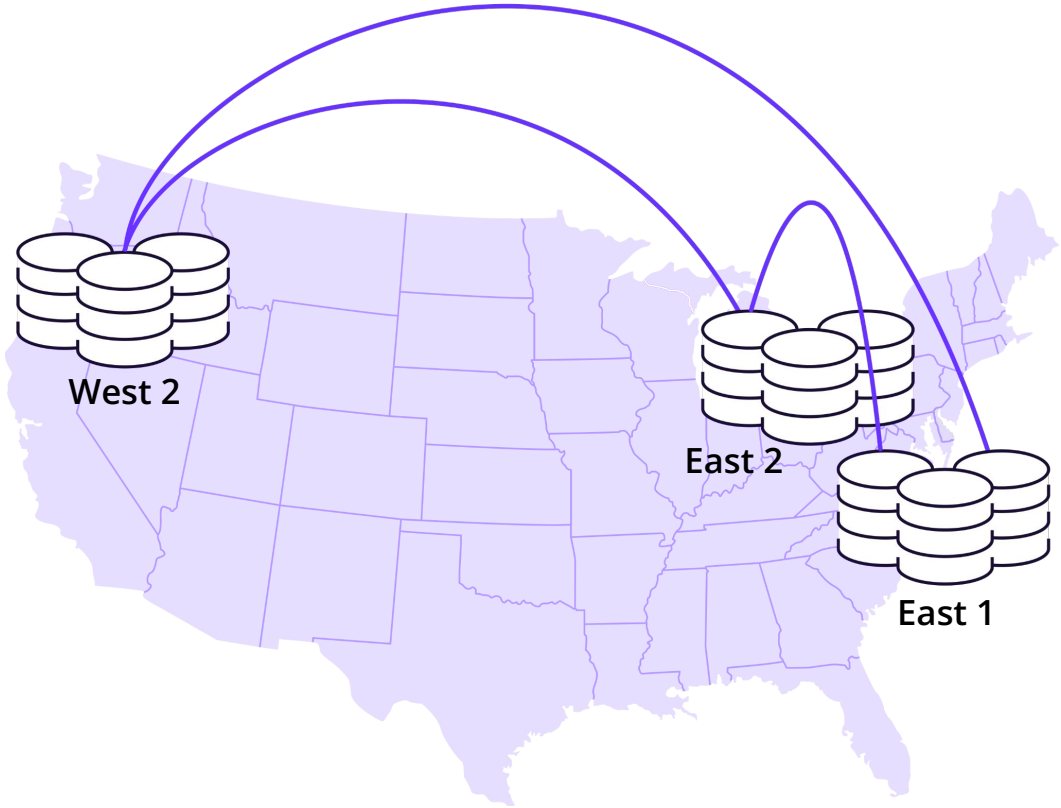

Running critical, distributed workloads with serializable isolation

Achieving zero downtime region resiliency with 3+ data centers

Resiliency powered by native Raft replication and distributed SQL

CockroachDB's native consensus-based Raft replication

- Full data consistency across nodes, zones & regions
- Zero data loss during failures within a cluster

Raft Replication

But what if you have other requirements?

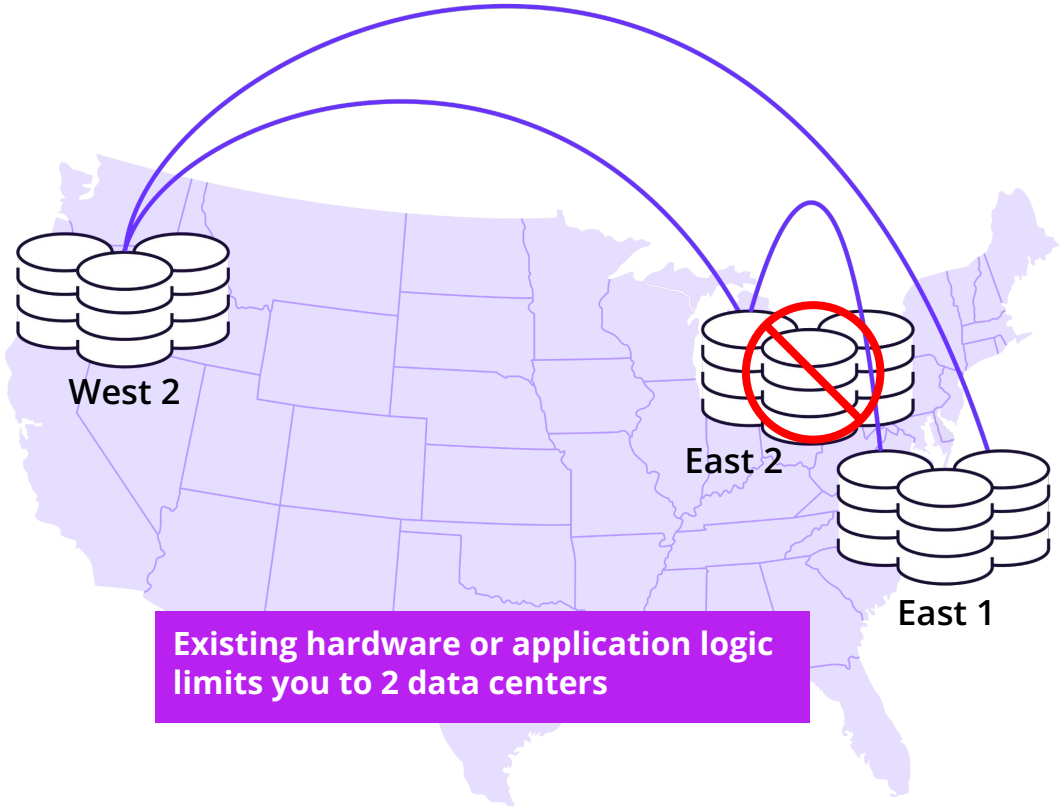

- Running critical, distributed workloads with serializable isolation
- Achieving zero downtime region resiliency with 3+ data centers
- You want to survive a region failure with 2 regions
- You want to achieve low, single-region write latency
- You want to migrate data between clusters
- You want more flexibility with deployment topologies

Introducing new tools to meet custom requirements

CockroachDB's native consensus-based Raft replication

- Full data consistency across nodes, zones & regions
- Zero data loss during failures within a cluster

Raft Replication

+

[GA] Physical Cluster Replication

- Transactionally consistent, full cluster replication
- Ensures data protection and disaster recovery with a passive standby cluster

+

[New] Logical Data Replication

- Eventually consistent, table and database level replication
- Enables high availability and real-time analysis with active clusters

Cross-Cluster Replication Tools

Physical Cluster Replication

Generally Available in 24.1

What is Physical Cluster Replication (PCR)?

Physical Cluster Replication continuously replicates all data from a *primary* CockroachDB cluster to an independent *standby* CockroachDB cluster

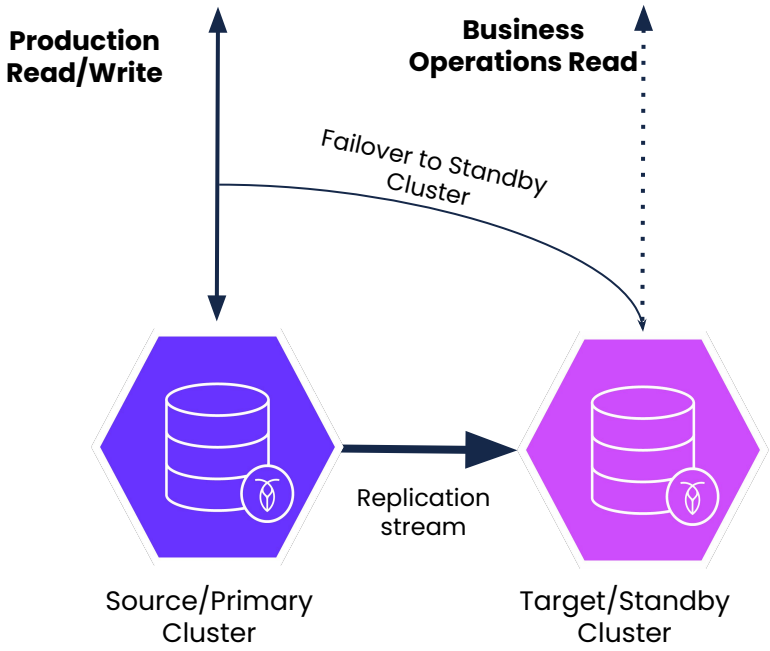

Use cases

Region and Cloud Survivability

Use PCR to survive an regional, cluster, or cloud outage with low RPO and RTO by cutting over to the standby.

Workload Isolation

Use PCR to offload non-critical operations, such as running analytics queries, to the standby cluster (24.3)

Data Migration

Use PCR to migrate data seamlessly between CockroachDB clusters.

- Cluster upgrade protection
- Populate dev clusters in real time

Physical Cluster Replication (PCR) overview

Cluster-level
Replication happens at cluster-level, i.e., it is unaware of database, table, row boundaries, etc.

Transactionally Consistent
You see a transactionally consistent state when you "failover" to the standby, including schema changes, zone configs, users, privileges, etc. No conflicts.

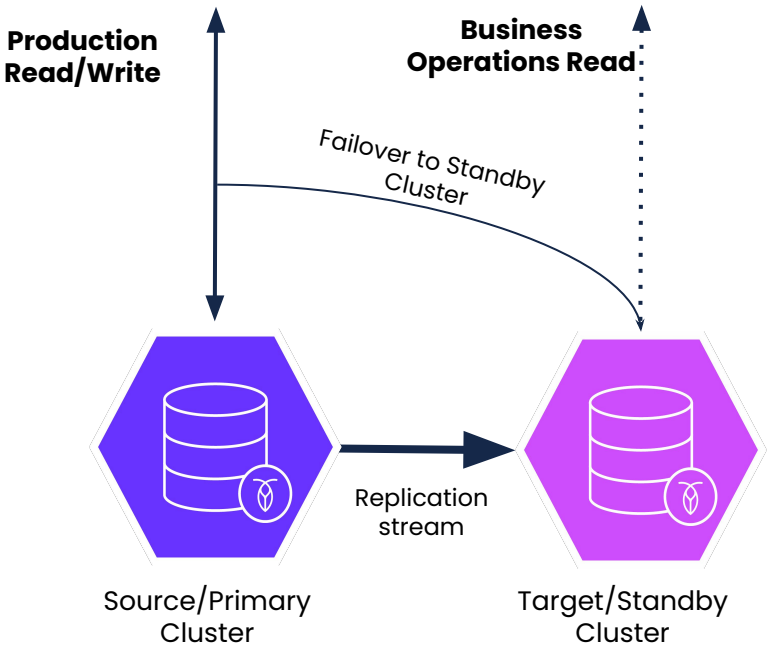

Life of a transaction in PCR

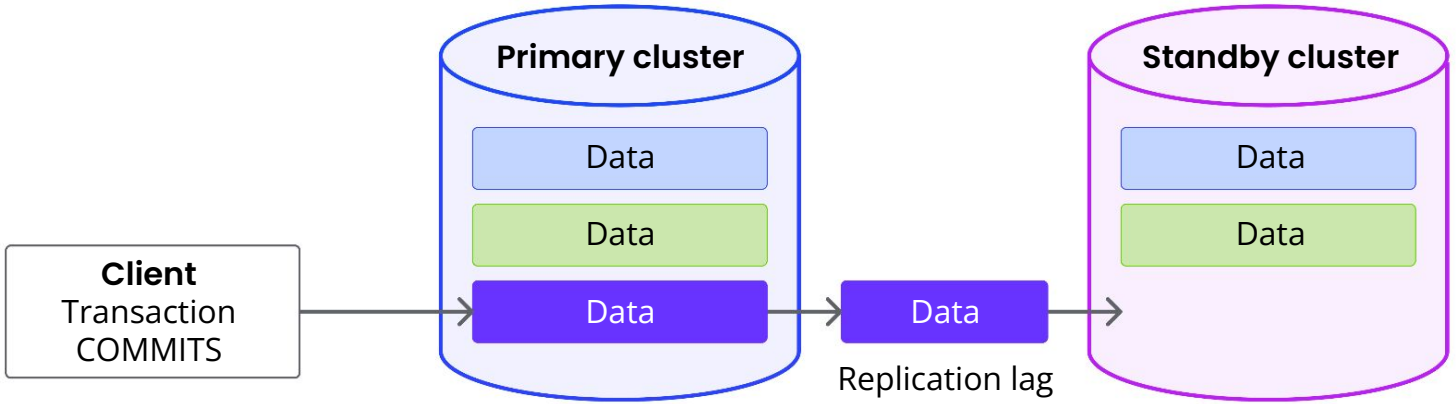

All data is accounted for in PCR during replication.

PCR replicates the entire cluster

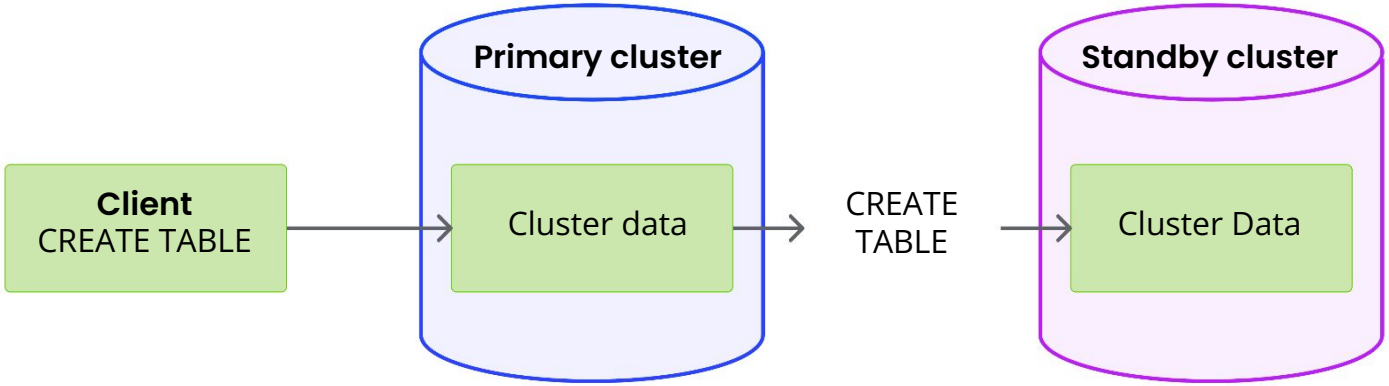

PCR replicates the entire cluster

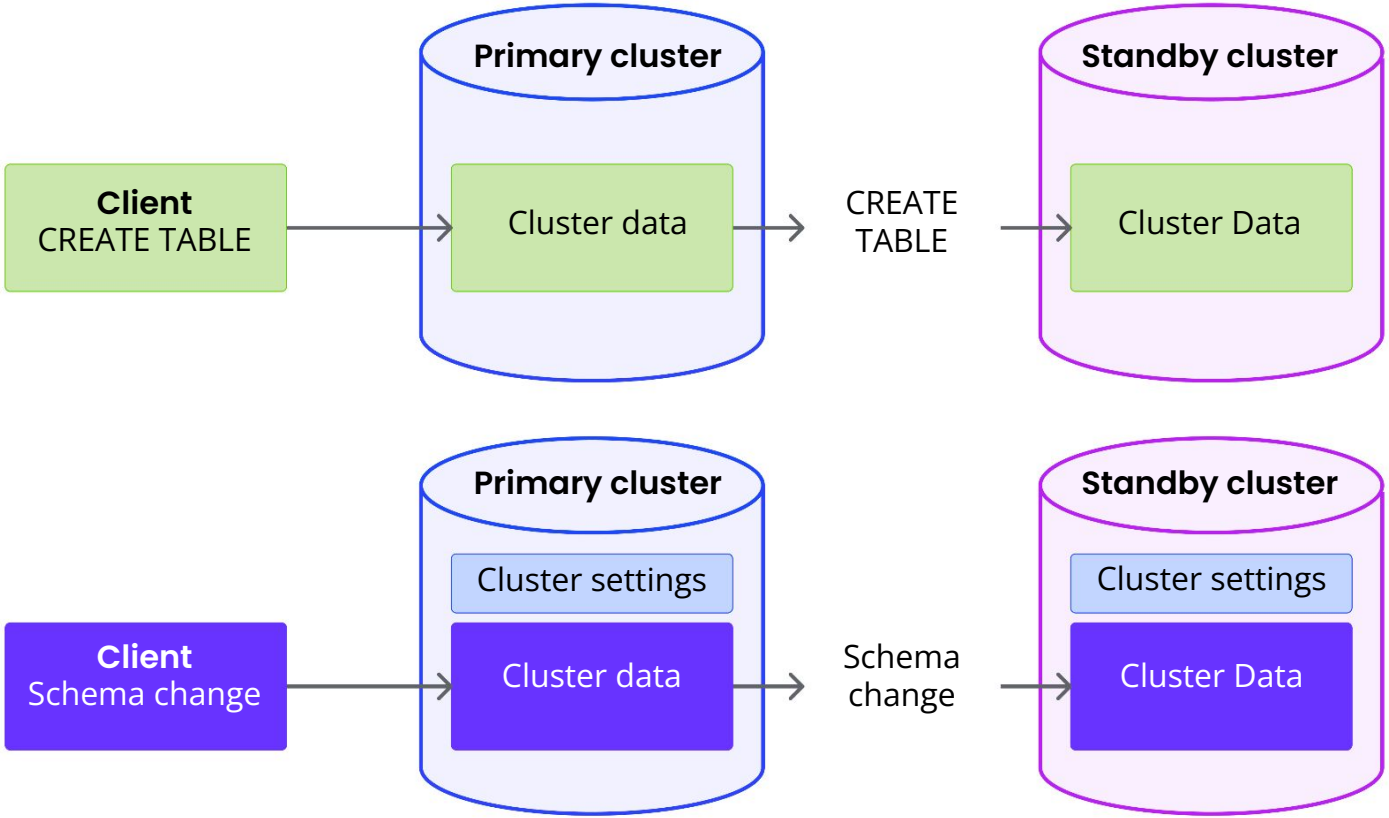

How does PCR handle disasters?

Node Outage, Network Partition

PCR will replan node connections between clusters and continue replication without any manual intervention required.

Cluster, Data-center, Region, Cloud Outage

Users can cutover to the standby cluster with a ~30 second RPO. Users will manually redirect application traffic to the standby cluster.

Human Error, dropping a table

PCR lets you cutover to a **time in the past,** ensuring that the target cluster is a consistent copy of the source cluster at the chosen timestamp.

Cyber-attack

PCR lets you cutover to a **time in the past,** ensuring that the target cluster is a consistent copy of the source cluster at the chosen timestamp.

Physical Cluster Replication (PCR) overview

RPO is ~30 seconds for most applications, depending on the workload traffic and corresponding replication lag.

RTO is ~one minute limited by the time it takes to identify the incident, decide and complete the cutover from the primary to standby database, & promote the standby to primary within your load balancer.

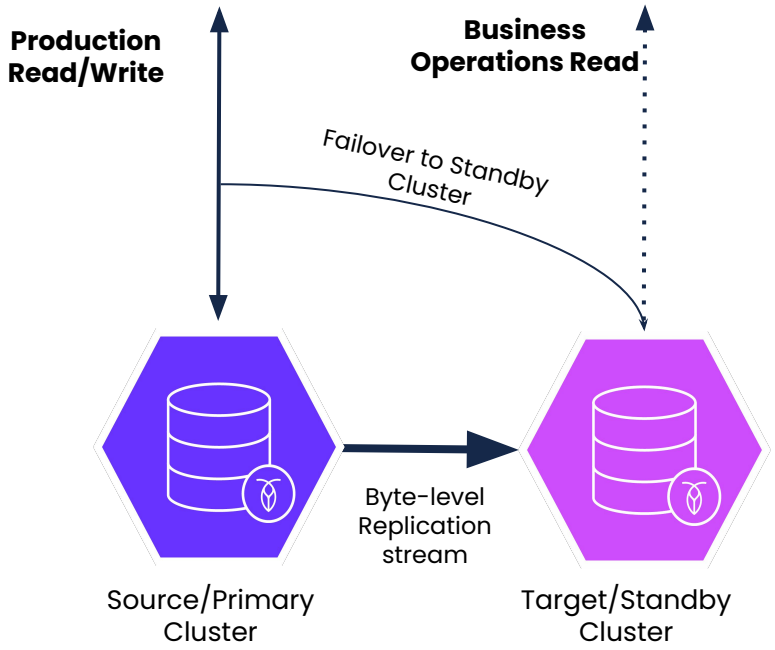

Physical Cluster Replication: roadmap

23.2	24.1 (May 2024)	24.3 (Nov 2024)	2025
Preview	Generally Available	Generally Available	Generally Available
Production ready	Full feature set supported for production use cases at scale (benchmarking at 100k writes/second) Fast cutback support	SELECT from standby **Preview on Cockroach Cloud**	BACKUP from standby **Generally Available on Cockroach Cloud**

Logical Data Replication

Public Preview in 24.3

What is Logical Data Replication (LDR)?

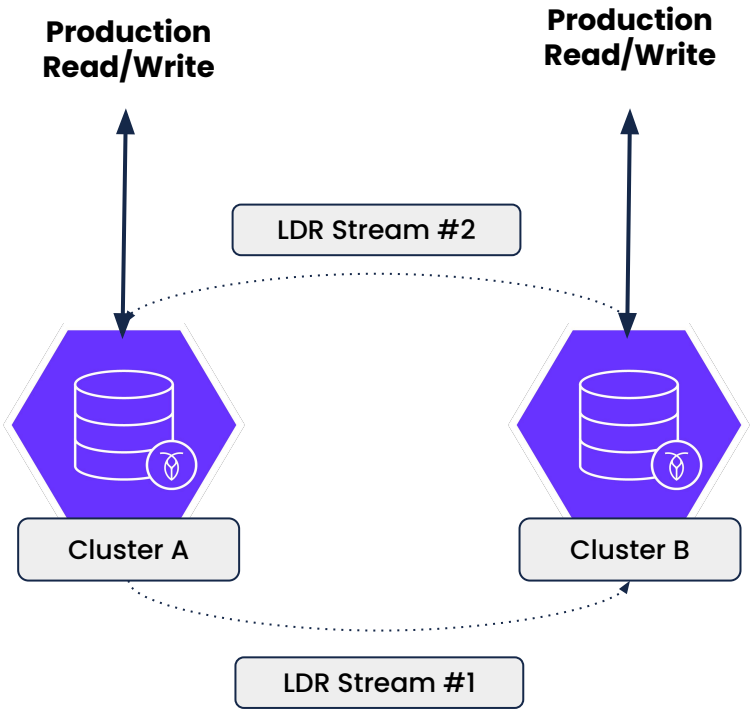

LDR provides **ACID consistency on each cluster** and **asynchronous replication** between clusters so that users can achieve **eventual consistency** between 2+ clusters.

Reads and writes can occur concurrently on both clusters allowing customers to survive a region failure while providing low latency read and writes.

What is Logical Data Replication (LDR)?

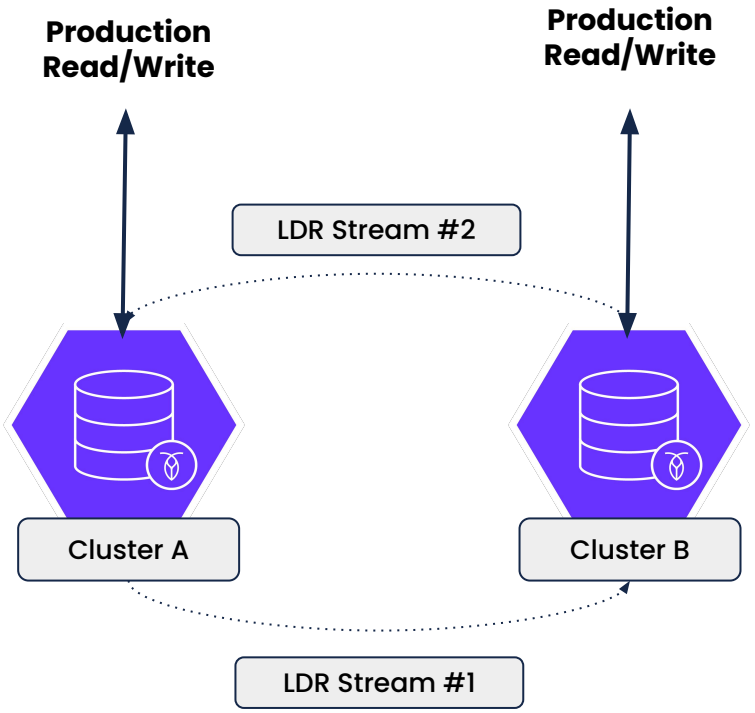

LDR replicates **table and database subsets** of a cluster.

LDR uses **last writer wins** conflict resolution based on the MVCC timestamps of the replicating write.

Use cases

Achieving High Availability with 2DC

Workload Isolation

Unlocking new deployment topologies

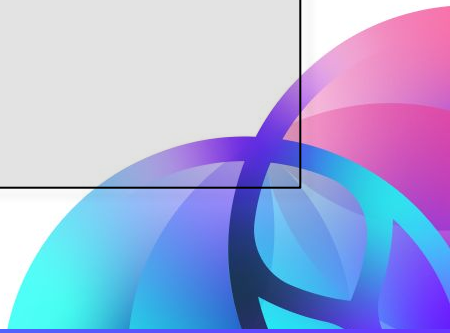

Achieving high availability and low latency with 2 data centers

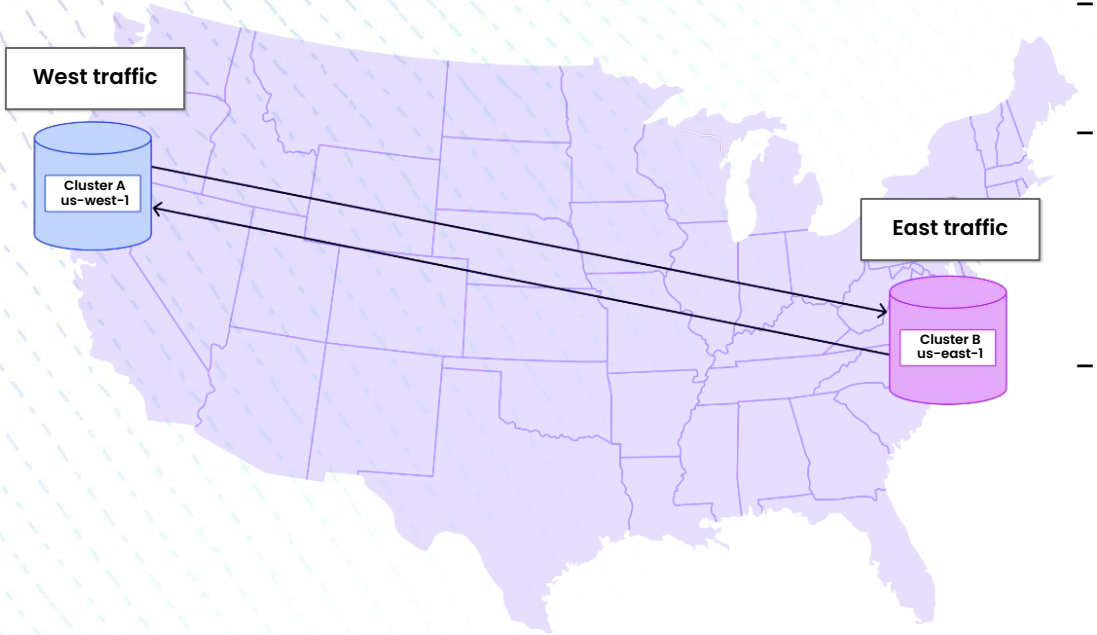

- Both clusters can receive application reads and writes
- Both clusters have low, single-region write latency with transactionally consistent writes via Raft consensus replication
- In a data center or cluster outage, operators can redirect application traffic to the surviving cluster with little downtime

Workload isolation

LDR enables users to isolate critical application workloads from non-critical application workloads, enabling users to isolate resources and size clusters for certain workloads.

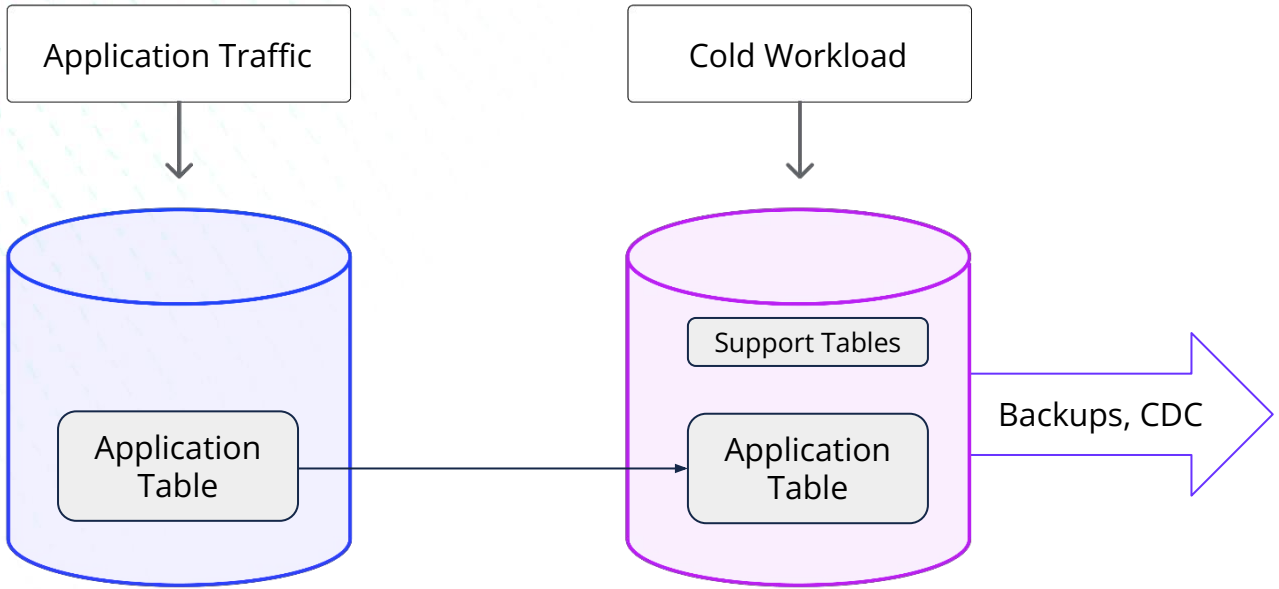

Unlocking new deployment topologies

Hub and Spoke

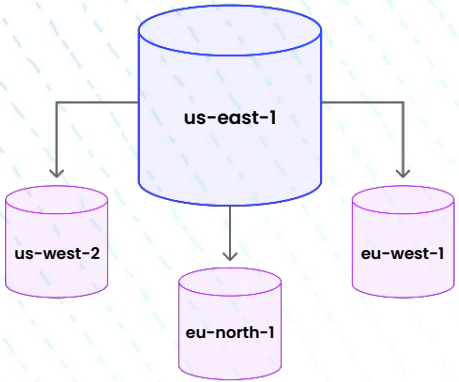

Achieve low read latencies with spoke clusters

Fan-In

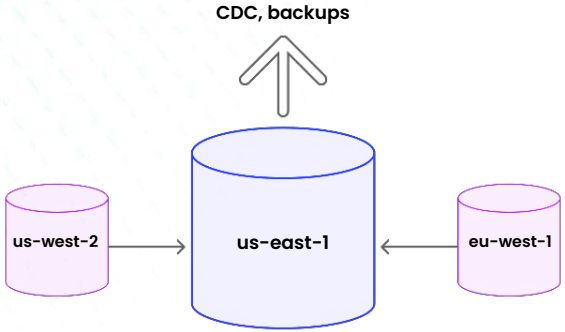

Offload non-critical operations to a sink cluster

Fully Connected

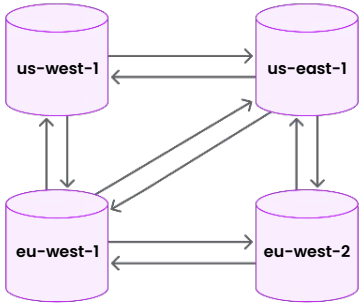

Achieve single-region read and write latencies while maintaining a highly available global system

Logical Data Replication: Roadmap

24.2	24.3 (Nov 2024)	2025	
Private Preview	Public Preview	Generally Available	
Testing with enterprise design partners	**Ready for production workloads**	Database level replication Improved schema change replication	Transactional mode Custom conflict resolution **Preview on Cockroach Cloud**

Cross-Cluster Replication in CockroachDB

	Physical Cluster Replication (PCR)	Logical Data Replication (LDR)
What is it?	• **One-way physical replication** • Replicates the entire source cluster	• **Bi-directional logical replication** • Replicate table or database subsets of a cluster
Why would you use it?	**Surviving region/DC outages with 2 DCs** **Surviving cluster control plane outages** **Data migration between clusters** **Create flexible deployment topologies**	
Operator Experience	**Started by SQL statements, in the DB** **Jobs infrastructure integration** **Admission Control for protecting foreground workloads**	

Thank you