



Modern Distributed Databases: Architecture, DevOps, and Operational Trade-offs

Srikanth Uddanti | Sr Manager, Software Development and Engineering | Charles Schwab Inc

The Challenge with Monolithic Databases

As cloud-native platforms scale, traditional monolithic databases hit hard architectural ceilings. Three failure modes consistently emerge at enterprise scale:

High Availability

Single points of failure undermine uptime SLAs

Global Distribution

Latency and consistency suffer across regions

Real-Time Processing

Throughput ceilings block event-driven workloads

Session Agenda

01

Foundational Concepts

CAP theorem, replication models, and sharding strategies

03

DevOps Automation

IaC, configuration management, and continuous deployment

02

Consensus Protocols

Raft and Paxos for fault tolerance and availability

04

Security and Governance

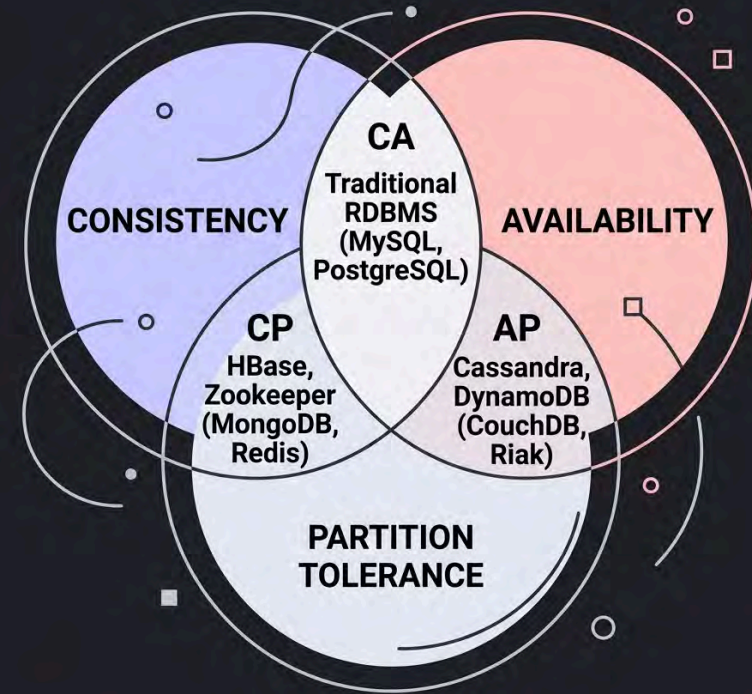
Zero-trust, certificate auth, and distributed key management

Understanding the CAP Theorem

The CAP theorem states that a distributed system can guarantee only two of the following three properties simultaneously:

- **Consistency:** Every read receives the most recent write
- **Availability:** Every request receives a response
- **Partition Tolerance:** System operates despite network partitions

i In practice, network partitions are unavoidable in distributed systems, forcing architects to choose between consistency and availability during failure scenarios.



Replication Models

Single-Leader

All writes go to one primary node; replicas serve reads. Simple but creates a bottleneck and potential single point of failure.

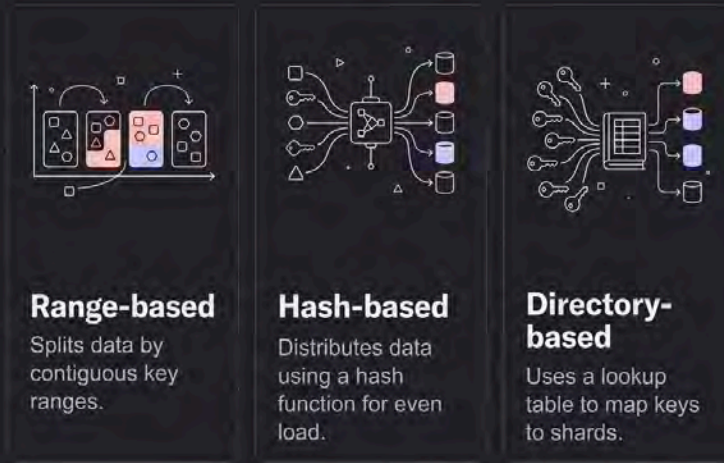
Multi-Leader

Multiple nodes accept writes simultaneously. Enables geo-distributed writes but introduces conflict resolution complexity.

Leaderless

Any node can accept reads and writes using quorum-based consistency. Maximizes availability but requires careful tuning of read/write quorums.

Sharding Strategies



Sharding horizontally partitions data across nodes to eliminate single-node capacity limits. Selecting the right strategy directly impacts query performance, hotspot risk, and rebalancing complexity.

→ Range-Based

Good for range scans; susceptible to hotspots on monotonic keys

→ Hash-Based

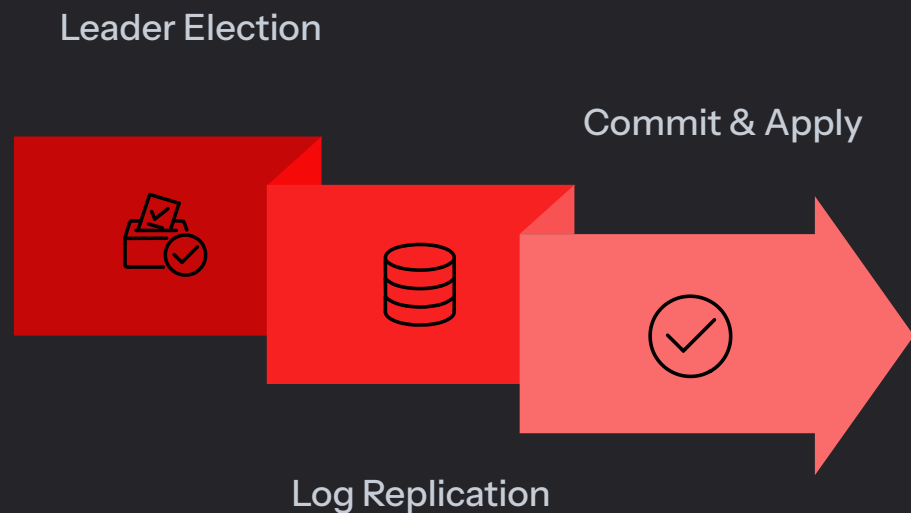
Uniform distribution; sacrifices range query efficiency

→ Directory-Based

Maximum flexibility; lookup table becomes a critical dependency

Raft: Understandable Consensus

Raft decomposes consensus into three relatively independent subproblems: leader election, log replication, and safety. A single elected leader handles all client writes, appending entries to its log and replicating them to followers before committing. This design makes Raft significantly easier to reason about and implement correctly compared to Paxos, which is why it underlies databases like CockroachDB, TiKV, and etcd.



Paxos vs. Raft at a Glance

Paxos

- Original consensus algorithm by Leslie Lamport
- Highly flexible and mathematically proven
- Notoriously difficult to implement correctly
- Multi-Paxos required for practical log replication
- Used in Google Spanner and Chubby

Raft

- Designed explicitly for understandability
- Strong leader model simplifies reasoning
- Formalized leader election via randomized timeouts
- Easier to implement, test, and verify
- Used in etcd, CockroachDB, TiDB, and TiKV

ⓘ Both protocols guarantee safety (no two nodes commit different values) but differ sharply in implementation complexity and operational observability.

Infrastructure as Code for Data Platforms

Treating database infrastructure as versioned, declarative code eliminates configuration drift, accelerates environment provisioning, and enables audit trails for every change.



Provisioning

Terraform and Pulumi codify cluster topology, node sizing, and network policies for repeatable deployments.



Configuration Management

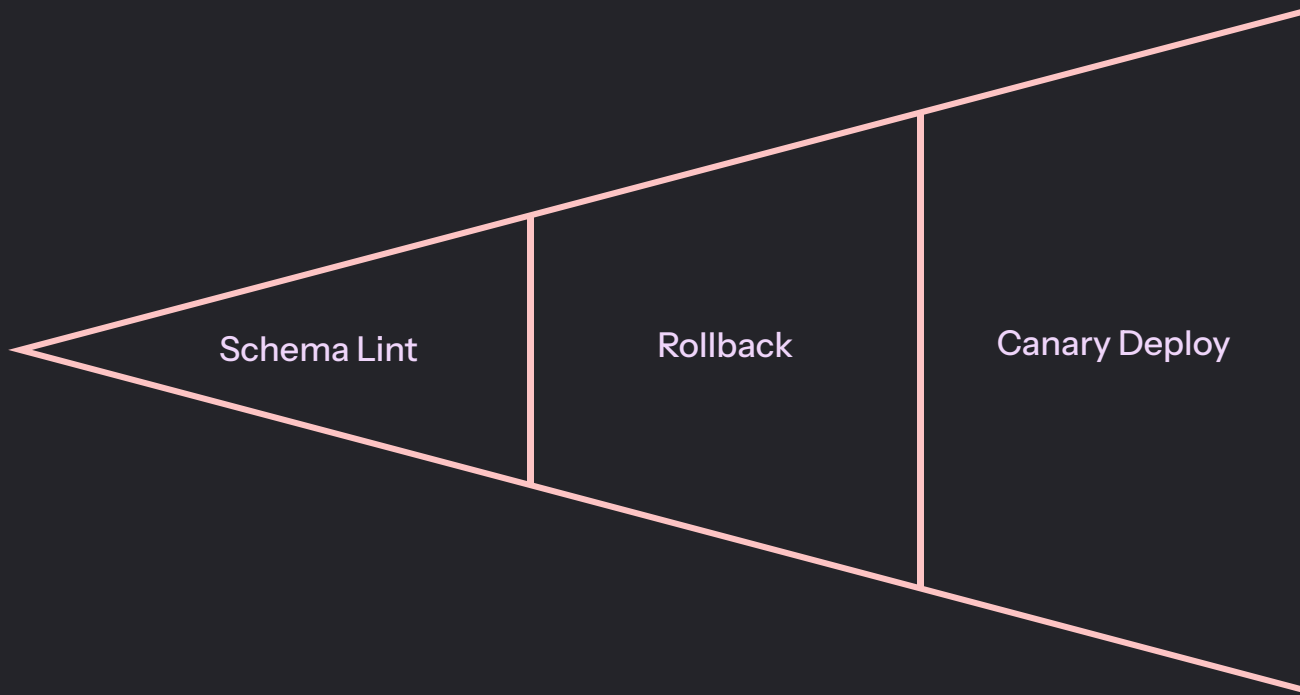
Ansible and Helm charts enforce consistent tuning parameters across dev, staging, and production clusters.



Continuous Deployment

GitOps pipelines automate schema migrations, rolling upgrades, and canary rollouts with automated rollback gates.

Continuous Deployment Patterns for Databases



Database deployments require additional caution beyond stateless services. Key automation primitives include:

- **Expand-Contract migrations** to avoid breaking backward compatibility
- **Blue-green clusters** for major version upgrades with instant cutover
- **Automated health gates** checking replication lag, query latency, and error rates before promotion

Reliability and Scalability Through Automation

99.99%

Target Uptime

Achievable with automated failover and self-healing cluster operators

~0

Manual Toil Goal

Kubernetes operators automate backups, failover, scaling, and cert rotation

3x

Faster Recovery

Automated runbooks reduce MTTR compared to manual incident response

Operator patterns (e.g., Vitess, CockroachDB Operator, MongoDB Atlas Operator) encode operational knowledge directly into the control plane, enabling self-healing behavior without human intervention.

Zero-Trust Infrastructure for Databases

Zero-trust assumes no implicit trust for any node, user, or service, whether inside or outside the network perimeter. Applied to distributed databases, this means:

Mutual TLS (mTLS)

All inter-node and client-to-node communication is encrypted and authenticated with certificates.

Identity-Based Access

Fine-grained RBAC policies enforced at the database layer, not just the network perimeter.

Audit Logging

Every query, schema change, and admin action is captured and shipped to a tamper-resistant log store.

Certificate-Based Auth and Distributed Key Management

Certificate Lifecycle

Automated certificate issuance and rotation via Vault PKI or cert-manager eliminates manual cert management and reduces expiry-related outages.

Distributed KMS

Encryption keys managed externally in HashiCorp Vault or cloud KMS (AWS KMS, GCP Cloud KMS) with automatic key rotation policies.

Secrets Management

Dynamic database credentials issued per-service with short TTLs prevent long-lived credential exposure and lateral movement.

Operational Trade-offs Summary

Every architectural decision in a distributed database involves deliberate trade-offs. Understanding them is the foundation of sound operational strategy.

Dimension	Strong Consistency	High Availability	Key Consideration
Replication	Synchronous	Asynchronous	RPO vs. write latency
Consensus	Raft / Paxos	Eventual consistency	Quorum size vs. throughput
Sharding	Directory-based	Hash-based	Flexibility vs. lookup overhead
Deployment	Blue-green cluster	Rolling upgrade	Cost vs. risk tolerance

Key Takeaways

Know your consistency model

CAP and PACELC trade-offs must align with business SLAs before any technology selection.

Automate everything operational

IaC, GitOps pipelines, and Kubernetes operators are not optional at enterprise scale; they are load-bearing infrastructure.

Build security into the data plane

Zero-trust, mTLS, and dynamic secrets must be designed in from day one, not retrofitted after a breach.

Thank you.