

CLUSTRIX 9

The only scale-out, drop-in MySQL replacement for high transaction, high-value workloads.

ClustrixDB 9

With ClustrixDB release 9, Clustrix continues moving forward in its mission to provide the leading scale-out SQL database that provides application developers with unparalleled scale, high-availability, ACID compliance and database-embedded logic for cleaner application code and data models. ClustrixDB accomplished this without sharding, replication, or master-slave configurations.

ClustrixDB 9 is for OLTP applications that need the scalability of a distributed database without sacrificing availability or data integrity and that process modern unstructured data and microsecond events – applications such as mobile, e-commerce, social, SaaS, gaming IoT and ad tech. With ClustrixDB 9, developers who want to innovate do not have to choose between an RDBMS that has these features but does not scale well and a NoSQL database that has to relax availability and consistency in order to scale.

Key capabilities of ClustrixDB 9 include:

- AWS Availability Zones deployments
- JSON Support for Semi-structured and Schema-less Data
- Fractional Record and Event Timestamps to Support Microsecond Event Data
- Generated Columns for Commonly Calculated Values

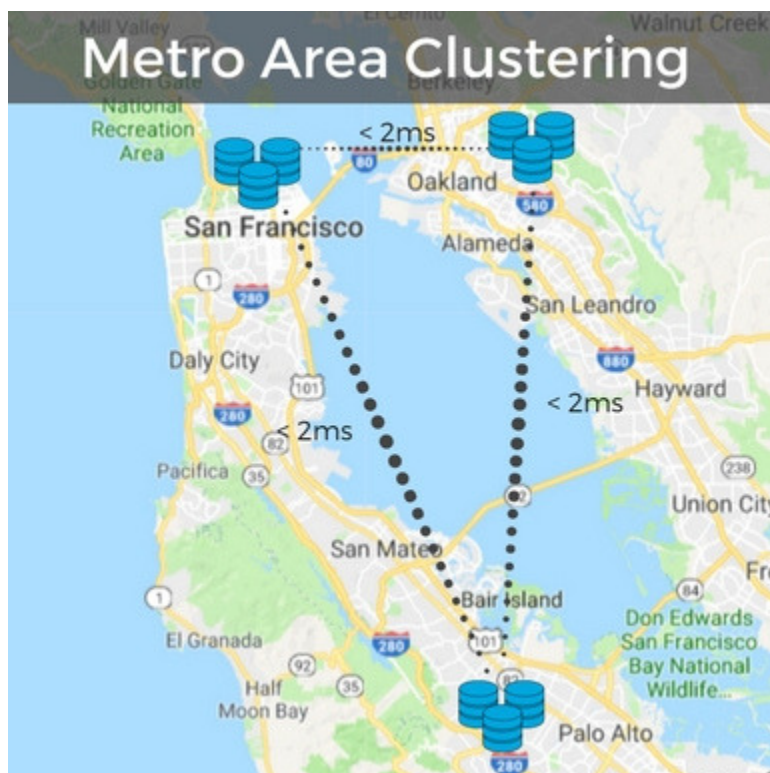
METRO AREA CLUSTERING AND AVAILABILITY ZONES

For the increasing number of applications moving to the cloud to take advantage of the lower TCO and elastic scalability, ClustrixDB 9 adds support for metro area clustering and availability zones, including AWS AZs, in public or private clouds. With this release, Clustrix continues its commitment to provide application developers with both scale and high-availability so they don't have to choose one over the other.

ClustrixDB 9 supports deploying a single OLTP database across multiple AWS Availability Zones (AZ) that is fault tolerant in the face of zone failures.

Availability and Data Protection

ClustrixDB will automatically keep the database online if there is a complete zone failure without the DBA or application having to do anything. Data will stay consistent across nodes, clusters and zones.





Data Integrity

ClustrixDB AWS AZ support enables each node to be a read-write node, regardless of which zone it's in so that data stays consistent across nodes, clusters and zones. This provides superior data consistency than with replication or master/slave configurations.

Scalability and Performance

ClustrixDB multi-AZ support allows more write-power with multiple read-write nodes in different zones without rewriting the application. This provides better performance and lower latency than with replication or single write-nodes in each zone.

The Clustrix Difference

Unlike other distributed RDBMS available on AWS, such as MySQL and its derivatives, ClustrixDB does not limit the number of read-writes nodes in each zone to a single master. And unlike NoSQL databases, ClustrixDB has the intelligence to handle multiple read-write nodes automatically, without embedding database-specific logic in the application.

MODERN DATA

ClustrixDB 9 can handle more sophisticated, modern data such as semi-structured data, fractional-second events, and generated columns while simultaneously delivering the advantages of the ClustrixDB distributed RDBMS:

- Scalability without sharding or replication
- High throughput of a NoSQL database without sacrificing ACID properties such as availability or consistency
- Reduced burden on application development
- Fast performance

JSON Support for Semi-structured and Schema-less Data

ClustrixDB 9 supports a native JSON data type and functions for parsing data and indexes to support JSON queries based on the content within JSON data or documents. JSON data can be joined with transactional data, while maintaining high scalability without replication or sharding, and full ACID compliance without relaxing availability or consistency.

Fractional Record and Event Timestamps to Support Microsecond Event Data

ClustrixDB 9 supports storing dates and times with precision up to the microsecond, enabling applications to maintain the critical timestamp granularity needed for event-based applications.

Generated Columns for Commonly Calculated Values

ClustrixDB 9 supports generated columns for calculated columns that can be virtual (calculated on the fly) or stored and indexed. Column values can be generated from other columns by defining expressions at the table level. Queries on generated columns automatically get the right value every time, changes to application logic can be reduced and commonly used values can be indexed for fast retrieval. Combined with ClustrixDB 9's JSON support, developers can dynamically create indexes within JSON documents for better search and query performance.

For a complete list of new features in Clustrix 9, request release notes from info@clustrix.com.

ClustrixDB

Flexible transactional scale for the connected world.

www.clustrix.com | info@clustrix.com | 877.806.5367