

Oracle NoSQL Database 22.3 Enterprise Edition (EE)

Oracle NoSQL Database is a multi-model, multi-region, multi-cloud, active-active database designed to provide a highly available, scalable, performant, flexible, and reliable data management solution to meet today's most demanding workloads. It can be deployed in enterprise data centers and cloud environments. It is well-suited for high volume and velocity workloads, like Internet of Things, 360-degree customer view, online contextual advertising, fraud detection, mobile application, user personalization, and online gaming. Developers can use a single application interface to build applications that run in enterprise and cloud environments.

Applications send network requests against an Oracle NoSQL data store to perform database operations. With multi-region tables, data can be globally distributed and automatically replicated in real-time across different regions. Data can be modeled as fixed-schema tables, documents, key-value pairs, and large objects. Different data models interoperate with each other through a single programming interface.

Oracle NoSQL Database is a sharded, shared-nothing system that distributes data uniformly across multiple shards in a NoSQL database cluster, based on the hashed value of the primary keys. An Oracle NoSQL Database data store is a collection of storage nodes that hosts one or more replication nodes. Data is automatically populated across these replication nodes by internal replication mechanisms to ensure high availability and rapid failover in a storage node failure.


Oracle NoSQL Database provides Java, Python, Node.js, Go, C#, C, Spring Data drivers and REST API for application development. It integrates with other Oracle products and open source applications to simplify and streamline application development. Oracle NoSQL Database is dual-licensed and available as an open-source Apache-licensed Community Edition as well as a commercially licensed Enterprise Edition.

KEY BUSINESS BENEFITS

- High throughput
- Bounded latency
- Linear scalability
- High availability
- Fast and easy deployment
- Smart topology management
- Online elastic configuration
- Multi-region data replication
- Enterprise grade software and support

Architecture

Oracle NoSQL Database is built upon the proven Oracle Berkeley DB Java Edition high-availability storage engine, widely used in enterprises across industries.



A layer of services is implemented on top of the storage engine to provide a distributed database architecture. Oracle NoSQL Database is designed for enterprise applications requiring predictable low latency, on-demand scalability, data model flexibility, effortless operations.

High Availability and No Single Point of Failure

Each shard in the Oracle NoSQL Database provides dynamically elected leader nodes (masters) and multi-replica database replication. Transactional data is delivered to all replica nodes in the shard with flexible durability policies per transaction. If the master replica node of a shard fails, a PAXOS-based automated failover election enables a new shard master with minimal effects to write latency. This allows for scalability, high availability, and low latency read and write operations.

High Performance

Oracle NoSQL Database is network topology and latency aware. The Oracle NoSQL Database Driver working in conjunction with a highly scalable, fault-tolerant, high throughput storage engine enables a more granular distribution of resources and processing, which reduces the incidence of hot spots and provides excellent performance on commodity-based hardware.

Transparent Load Balancing

The Oracle NoSQL Database Driver partitions the data in real time and evenly distributes it across the storage nodes. It is network topology and latency-aware, routing read and write operations to the most appropriate storage node to optimize load distribution and performance.

Configurable Smart Topology

Administrators can specify the capacity available on a given storage node, allowing storage nodes with higher capacity to host multiple replication nodes. Once the NoSQL cluster knows about the capacity for the storage nodes in a configuration, it automatically allocates replication nodes intelligently. This

KEY FEATURES

- Multi-region Table
- KVLocal for edge computing
- Single common application programming interface
- Native JSON data type
- JSON data type queries with SQL
- Java, Python, Node.js, Go, C, C#, Spring Data Drivers
- Fast, index Btree storage
- Dynamic partitioning (sharding)
- Transparent load balancing
- Streaming Large Object support
- Multi data models
- GeoJSON support
- Secondary index support
- Streams support
- ACID compliant transactions
- Replication for HA, fault tolerance, failover, read scalability
- JMX for system monitoring
- Online rolling upgrade
- Easy data migration
- Efficient multi-zone support
- Data encryption using SSL
- Node level backup and restore
- Integrated with Apache Hadoop
- Secure full text search
- Aggregation
- Parent child joins
- Zone affinity
- Third party development tool and framework support

results in better load balancing for the system, better use of system resources, minimizing system impact in the event of any storage node failure.

Elastic Configuration

Oracle NoSQL Database includes a topology planning feature, with which an administrator can now modify the configuration of a NoSQL database while the database is still online. This allows administrators to:

- **Increase Data Distribution** by adding more shards in the cluster to achieve better write throughput.
- **Increase Replication Factor** by assigning additional replication nodes to each shard, increasing read throughput and system availability.
- **Rebalance Data Store** by modifying the capacity of a storage node(s). The system can be rebalanced, re-allocating replication nodes to the available storage nodes, as appropriate. The topology rebalance command allows administrators to move replication nodes and/or partitions from overutilized nodes onto underutilized storage nodes or vice versa.

Easy Administration and Enhanced System Monitoring

Oracle NoSQL Database provides an administration service, which can be accessed from the command-line interface (CLI). This service supports core functionality such as configuring, starting, stopping, and monitoring a storage node without requiring manual effort with configuration files, shell scripts, or explicit database operations.

Administrators can deploy the Java Management Extensions (JMX) agents to monitor the cluster health. The agents poll information about the status, performance metrics, and operational parameters of the storage node and its managed services.

Arbiters

An arbiter is deployed to reduce the hardware resource requirement in a cluster, achieving fewer replicas per shard. It does not store a copy of the data and cannot become a master but participate in elections to select a master.

Online Rolling Upgrade

Administrators can perform rolling upgrades to all nodes in the Oracle NoSQL Database cluster while the database remains online and available to client applications.

Multi-Region Table

Oracle NoSQL Database multi-region, active-active architecture enables the creation of Multi-Region Tables in geographically distributed NoSQL clusters, maintaining consistent data across. A Multi-Region Table is a global, logical table that is deployed and maintained in different regions. It is a read-anywhere and

USE CASES

- “Last mile” Big Data connectivity
- Click-through data capture
- High-throughput event processing
- Fraud detection
- Metadata storage
- Social Network data capture
- Online retail customer view
- Mobile application back end infrastructure
- Real-time sensor aggregation
- Network device monitoring and management
- Scalable authentication
- Content management
- Archiving

write-anywhere table that lives in multiple regions. All Multi-Region Tables deployed in different regions are synchronized via NoSQL Streams. Each region's NoSQL cluster may have a different topology, different hardware configuration, different operations, and security management. Multi-Region Tables support CRDT in JSON document and fixed-schema data models for conflict-free data replication across regions.

Multi-Zone Deployment

Oracle NoSQL Database supports the definition of multiple zones from within the topology deployment planner. It leverages the definition of these zones internally to intelligently allocate replication of processes and data, ensuring optimal reliability during hardware, network, and power-related failures.

There are two types of zones. Primary zones contain master or replica nodes connected by fast interconnects, serving write and read operations. These zones participate in master elections and commit acknowledgments. Secondary zones have replicas nodes only, providing low latency read access to data at a distant location or maintaining an extra copy of the data to increase redundancy. These zones can offload read-only workloads from primary zones for analytics, report generation, improving workload management.

Oracle NoSQL Database enables users to continue business operations in the event of zone failures. It also allows planned maintenance by taking one or more zones offline without impacting business operations. Additionally, with the zone affinity feature, it's possible to place master nodes in primary zones that are in close network proximity to the user applications, ensuring predictable write latencies.

KVLocal

KVLocal is an embedded NoSQL database which can be viewed as a single-node Oracle NoSQL Database that is robust and requires minimal administration. It runs as a separate child process in the application JVM. KVLocal can be configured to use either TCP/IP sockets or Unix domain sockets for communication between the client APIs and KVLocal. It is suited for edge computing.

Single Application Programming Interface

HTTP proxy is a middle-tier component between the client applications and the NoSQL Database server. The HTTP protocols are identical for Oracle NoSQL Database enterprise edition and cloud service, allowing client applications to easily connect and move between both products. With a common application programming interface, developers can easily build applications that run and interoperate between the NoSQL enterprise edition and cloud service.

Table Data Model

A NoSQL table is a collection of record items with the same record type in a fixed schema data model, resembling the relational tables in RDBMS. Records are layered on top of the distributed key-value structure, inheriting all its advantages

RELATED PRODUCTS

The following Oracle products can be used in conjunction with Oracle NoSQL Database:

- Oracle Big Data SQL
- Oracle Berkeley DB
- Oracle Coherence

and simplifying application design by enabling seamless integration with familiar SQL-based applications.

Native JSON Data Type

JSON data type stores schemaless data that does not conform to a rigid structure. Oracle NoSQL Database stores JSON documents in an internal binary (optimized) format that allows quick read access. Indexes can be created on any JSON object, including the deeply nested objects, to efficiently search and access the document.

Secondary Indexes

Secondary indexes offer an alternative way of retrieving table rows besides using primary keys. By creating a secondary index, queries can retrieve rows with different primary keys but share other characteristics. Oracle NoSQL Database supports rich secondary indexing capabilities, including indexing deeply nested arrays and maps, dramatically improve performance for queries. Indexes can be created for the value of the NoSQL built-in functions on timestamps, strings, and row properties.

SQL for NoSQL

Developers can use the SQL-like interface to perform CRUD operations from a command-line interface, scripts, or SDKs. The interface supports flat relational data, hierarchical typed (schema-full) data, and schemaless JSON documents, handling multiple data models seamlessly. The SQL Path Expression enables navigating inside complex data structures and selecting their nested values using different step operations.

Partial JSON Document Update

Developers can update (change, add, remove) a part of the JSON document. Such an update happens on the server side, eliminating the need for the read-modify-write cycle, which is atomic and thread-safe.

GeoJSON Support

Data can be stored in GeoJSON format to represent geographical features, properties, and boundaries. Geometry types supported are Point, LineString, Polygon, MultiPoint, MultiLineString, MultiPolygon, and GeometryCollection. Search functions support queries on geographical data that have a relationship based on a specific geometry. Indexes can be created for optimal search performance.

Hierarchical Tables

The NoSQL table hierarchy organizes tables in a parent-child relationship for storage efficiency in write-heavy workloads and fine-grained table access for security. It also enables joining the parent-child tables using the NESTED TABLES clause and LEFT OUTER JOIN clause. LEFT OUTER JOIN closely

resembles the semantics used in relational database queries. Oracle NoSQL Database supports table hierarchy deployment in Multi-Region tables, ushering in a new application development paradigm.

Aggregation Functions

Aggregate functions iterate over the data, evaluate an expression and aggregate the returned values into a single value. Supported functions are sum, count, avg, min, max.

Simple and Easy to Use APIs in Multiple Programming Languages

Oracle NoSQL Database provides Java, Python, Node.js, Go, C, C#, and Spring Data drivers that allow developers to choose the right interface for their application development.

Full-Text Search

Full-Text Search identifies natural-language documents that satisfy a query and optionally sorts the result by relevance. Oracle NoSQL Database integrates with the Elasticsearch engine to enable the full-text search capability.

Streams Processing

Streams processing offers a notification service that enables users to subscribe to all logical changes (e.g., table row puts and deletes) in Oracle NoSQL Database. Applications can be alerted based on these changes, which allow for asynchronous monitoring of data changes.

Time-To-Live

Time-To-Live (TTL) is a mechanism that allows applications to keep data for a specific period of time and automatically remove it when it expires. TTL is expressed as the amount of time data is allowed to live in the store. Data reaching its expiration timeout value can no longer be retrieved and will not appear in any store statistics.

Easy Data Migration

Oracle NoSQL Database Migrator offers an effortless way to migrate data between Oracle NoSQL Databases or other products such as OCI Object Storage and third-party NoSQL databases. Database Migrator exports data from the selected source and imports data into the sink, supporting documents and tables migrations. Users can transform the data by specifying one or more rules during the migration. The migrator also supports other source formats like AWS DynamoDB JSON, MongoDB JSON, CSV, and Apache Parquet.

Analytics Integrator

Oracle NoSQL Database Analytics Integrator copies data stored in a NoSQL Database Cloud Service table to a table created in Oracle Autonomous Data

Warehouse (ADW) in which analytics can be performed. With data stored in ADW, analytics can also be performed in Oracle Analytics Cloud or Oracle Analytics Server. Oracle NoSQL Database Analytics Integrator automatically copies data to OCI Object Storage and later uploads it to ADW.

Integration with Oracle Database External Tables

Oracle NoSQL Database can be accessed using Oracle Database's External Tables feature. Users' data remains in Oracle NoSQL Database. This capability enables applications to access data in external sources as if it were in a table in Oracle Database.

Integration with Apache Hive

Oracle NoSQL Database provides a set of interfaces and classes that satisfy the Hive programming model, which can be used to query data stored in Oracle NoSQL Database. Such integration enables running Hadoop MapReduce jobs against data stored in Oracle NoSQL Database.

Integration with Other Oracle Products

The integration of Oracle Golden Gate enables the transactional data to be replicated to a target Oracle NoSQL Database using Oracle NoSQL Handler. The Oracle NoSQL Handler streams change data capture into Oracle NoSQL Database using the Table API.

The integration of Oracle Coherence allows Oracle NoSQL Database to be used as a cache for Oracle Coherence applications. It allows applications to directly access cached data from Oracle NoSQL Database.

Large Object Support

Oracle NoSQL Database provides an interface you can use to read and write Large Objects (LOBs) such as audio and video files. A LOB is stored as a sequence of chunks whose sizes are optimized for the underlying storage system. The system chooses individual chunk sizes based on its knowledge of the underlying storage architecture and hardware.

Enterprise Security

OS-independent, cluster-wide password-based user authentication and Oracle Wallet integration enable greater protection from unauthorized access to sensitive data. Session-level Secure Sockets Layer (SSL) encryption and network port restrictions deliver greater protection from network intrusion. Oracle NoSQL Database also supports integration with Kerberos as an external authentication service, allowing easy integration with customers' existing applications which Kerberos already protects.

Third Party Development Tool and Framework Support

Developers can choose the tools and framework that they are familiar with to speed up the application development. For IDE, Oracle NoSQL Database is

integrated with IntelliJ and Eclipse to ease access to NoSQL datastore within the development environment. For data framework, Oracle NoSQL Database is integrated with Spring Data so that developers can focus on implementing the business logic instead of database access code. It offers a data abstraction layer that allows developers to migrate applications effortlessly from other databases to Oracle NoSQL Database. For source-code editor, Oracle NoSQL Database is integrated with Visual Studio Code.

Commercial Grade Software and Support

Oracle NoSQL Database is a commercial product fully supported by Oracle, giving organizations confidence when deploying Oracle NoSQL Database in the production environments.

Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2020, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Disclaimer: This document is for informational purposes. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document may change and remains at the sole discretion of Oracle Corporation.