

# Oracle NoSQL Database 21.1 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 on-premise data centers and cloud. 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 quickly build applications that run in on-premise 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 which 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, each of which 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 the event of a storage node failure.


Oracle NoSQL Database provides Java, Python, Node.js, Go, C#, and C drivers and REST API for application development. It is integrated 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, which is 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. In the event the master replica node of a shard fails, a PAXOS-based automated failover election enables a new shard master with minimal affects 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 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 greater 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 in order to optimize load distribution and performance.

### Configurable Smart Topology

System administrators indicate how much capacity is available on a given storage node, allowing storage nodes with higher capacity to host multiple replication nodes. Once the system knows about the capacity for the storage

### KEY FEATURES

- Multi-region Table
- Single common application programming interface
- Native JSON data type
- JSON data type queries with SQL
- Java, Python, Node.js, Go, C, C# 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, fail-over, 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

nodes in a configuration, it automatically allocates replication nodes intelligently. This results better load balancing for the system, better use of system resources and minimizing system impact in the event of any storage node failure. Smart Topology also supports Data Centers, ensuring that a full set of replicas is initially allocated to each data center.

## 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 the administrator to:

- **Increase Data Distribution:** by increasing number of shards in the cluster, this increases the write throughput.
- **Increase Replication Factor:** by assigning additional replication nodes to each shard, which increases 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 the administrator to move replication nodes and/or partitions from over utilized 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 command-line interface (CLI). This service supports core functionality such as the ability to configure, start, stop and monitor a storage node, without requiring manual effort with configuration files, shell scripts, or explicit database operations.

In addition, it also allows Java Management Extensions (JMX) agents to be available for monitoring. This allows management clients to poll information about the status, performance metrics and operational parameters of the storage node and its managed services.

## Arbiters

An arbiter can be deployed to achieve fewer replicas per shard. It does not store a copy of the data and cannot become a master. An arbiter can participate in elections to select a master. Arbiters are commonly deployed to reduce the hardware resource requirement in a cluster.

## Online Rolling Upgrade

Upgrades and patches are important aspects of any software support cycle. Oracle NoSQL Database provides utilities to perform a rolling upgrade, allowing a system administrator to upgrade all of the nodes in the Oracle NoSQL Database cluster while the database continues to remain online and available to client applications.

## 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

## Multi-Region Table

Oracle NoSQL Database multi-region, active-active architecture enables the creation of Multi-Region Tables in multiple NoSQL clusters and maintains consistent data across these clusters. The multi-region deployment may be geographically distributed. A Multi-Region Table is a global, logical table that is deployed and maintained in different regions. It is a read-anywhere and write-anywhere table that lives in multiple regions. All Multi-Region Tables deployed in different regions are synchronized via NoSQL Streams. Each region NoSQL cluster may have different topology, different hardware configuration, different operations and security management. Multi-Region Tables supports CRDT 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 & power related failures.

There are two types of Zones: Primary zones contain nodes that can be served as masters or replica nodes and are typically connected by fast interconnects. Secondary zones contain nodes which can only be served as replicas. Secondary zones can be used to provide low latency read access to data at a distant location, or to offload read-only workloads, like analytics, report generation, and data exchange for improved workload management. Oracle NoSQL Database allows users to continue business operations in the event of zone failures. This allows for any planned maintenance that results in the taking of 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. This ensures predictable write latencies.

## Single Application Programming Interface

HTTP proxy is a middle tier component that sits between the client applications and the NoSQL Database server. The HTTP protocols are identical for Oracle NoSQL Database on-premise and Oracle NoSQL Database Cloud Service. Client applications can connect and move between both products easily. With a single common application programming interface, developers can easily build applications that run and interoperate between on-premise and cloud.

## Table Data Model

A tabular data structure is available, which simplifies application data modeling by leveraging existing schema design core concepts. Table model is layered on top of the distributed key-value structure, inheriting all its advantages and simplifying application design even further by enabling seamless integration with familiar SQL-based applications.

## Native JSON Data Type

### RELATED PRODUCTS

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

- Oracle Big Data Appliance
- Oracle Exadata
- Oracle Big Data SQL
- Oracle Berkeley DB
- Oracle SQL Developer
- Oracle Spatial and Graph

JSON data type makes it easy to store data that doesn't conform to rigid schema. Only valid JSON documents can be stored providing automatic JSON document validation. JSON documents stored in JSON columns are converted to an internal binary (optimized) format that allows quick read access to document elements. JSON documents can be indexed to enable quick access to the nested attributes embedded deep within the document. Query JSON documents with the familiar SQL dialect. This powerful feature gives developers the ability to use SQL to query schemaless data.

## Secondary Index

Primary key based indexing limits number of low latency access paths. Sometime application needs a few non-primary-key based paths to support the workloads for the real-time system. Secondary index on any value field dramatically improves performance for queries. Oracle NoSQL Database supports rich indexing capabilities, including indexing of deeply nested arrays and maps.

## SQL for NoSQL

Oracle NoSQL Database provides a SQL-like interface that can be used from a command line interface, scripts, or SDKs. Such SQL interface supports flat relational data, hierarchical typed (schema-full) data, and schemaless JSON document. The SQL interface is designed to handle these data models seamlessly. The SQL Path Expression enables navigating inside complex data structure and selecting their nested values using different types of step operations.

## Partial JSON Document Update

Developers can update (change, add, remove) a part of JSON document. Such update happens on the server side eliminating the need for read-modify-write cycle, 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 certain geometry. Indexes can be created for optimal search performance.

## Hierarchical Tables

Oracle NoSQL Database supports hierarchical tables designed for write-heavy workloads and fine-grained table access authorization. It also supports joins among different tables that have the same hierarchical relationship between them. NESTED TABLES clause and LEFT OUTER JOIN are ways different tables

can be joined. Both are semantically equivalent. LEFT OUTER JOIN closely resembles the semantic used in relational database queries.

## **Aggregation Functions**

Aggregate functions in Oracle NoSQL Database iterate over the rows, evaluate an expression for each row, and aggregate the returned values into a single value. Syntactically, Aggregate functions appear in the SELECT clause. Supported Aggregate 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, and C# APIs that allow developers to choose the right driver for their application development.

## **Full Text Search**

Full Text Search provides the capability to identify natural-language documents that satisfy a query, and optionally to sort them by relevance. Oracle NoSQL Database integrates with Elasticsearch to enable full-text searching capability.

## **Streams Processing**

Streams processing offers a notification service that permits a user to subscribe to all logical changes (e.g., table row puts and deletes) made to 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 expired. TTL is expressed as the amount of time data is allowed to live in the store. Data which has reached 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 from other products. With its support for multi data models, the Migrator can ingest data modeled as fixed-schema tables, documents, key-value pairs, and large objects.

## **Oracle Database Integration via 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.

## Oracle Big Data SQL and Hive Integration

Oracle Big Data SQL is a common SQL access layer to data stored in Hadoop, HDFS, Hive and Oracle NoSQL Database. This allows customers to run queries on Oracle NoSQL Database from Hive or Oracle Database. Users can also run MapReduce jobs against data stored in Oracle NoSQL Database that's configured for secure access. It supports both primitive and complex data types.

## Integration with Other Oracle Products

The integration of Oracle SQL Developer enables viewing (Read-only) the data stored in Oracle NoSQL Database.

The integration of Oracle Event Processing provides read access to Oracle NoSQL Database via the Oracle NoSQL Database cartridge. Once the cartridge is configured, CQL queries can be used to query the data.

The integration of Oracle Semantic Graph is enabled through a Jena Adapter for Oracle NoSQL Database to store large volumes of RDF data (as triplets/quadruplets). This adapter enables fast access to graph data stored in Oracle NoSQL Database via SPARQL queries.

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 supports Large Object (LOB), such as audio and video files. These are stored as a sequence of chunks where individual chunk sizes are chosen automatically by the system based upon its knowledge of the underlying storage architecture and hardware. Such intelligent processing and storing of LOB permit low latency operations across mixed workloads with values of varying sizes.

## Apache Hadoop Integration

This integration supports reading data from Oracle NoSQL Database and prepares it for writing to a Hadoop system. To move data from Hadoop to Oracle NoSQL Database, data is read from the Hadoop system using the standard mechanisms and written to Oracle NoSQL Database using the APIs.

## 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. This allows easy integration with customers' existing applications which are already protected by Kerberos.



## Commercial Grade Software and Support

Oracle NoSQL Database overcomes a significant limitation faced by many enterprises considering the implementation of NoSQL databases — the need for full supportability. Oracle NoSQL Database is a commercial product fully supported by Oracle. This gives organizations the confidence and reduces the risk when deploying Oracle NoSQL Database in the production environments.

## 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 provide ease of access to NoSQL data store 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.

---

### 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](https://blogs.oracle.com)

 [facebook.com/oracle](https://facebook.com/oracle)

 [twitter.com/oracle](https://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.