

Oracle® NoSQL Database

Concepts Guide



Release 20.2

E85371-11

August 2020

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2011, 2020, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Conventions Used in This Book

v

1 Introduction to Oracle NoSQL Database

NoSQL Database Server Licensing	1-2
NoSQL Database Client Licensing	1-2
NoSQL Database Option Differences	1-2
Architecture	1-3
Replication Nodes and Shards	1-5
Replication Factor	1-5
Partitions	1-6
Zones	1-6
Arbiter Nodes	1-7
Topologies	1-7
Multi-Region Architecture	1-8
Cross Region Service	1-10
Life Cycle of Multi-Region Tables	1-11
Data Models	1-17
Consistency	1-18
Durability	1-18
Quorum	1-19
Administration	1-19
KVLite	1-20
Saving Admin CLI History	1-20
Monitoring	1-21
Troubleshooting	1-21
Access and Security	1-21
Integration	1-22
Oracle Database Mobile Server Integration	1-22
Hadoop Integration	1-22
Property Graph Integration	1-23
Oracle External Tables Integration	1-23

Preface

This document introduces Oracle NoSQL Database.

This book is aimed at technical users, primarily database administrators and developers who are new to Oracle NoSQL Database.

Conventions Used in This Book

The following typographical conventions are used within this manual:

Information that you are to type literally is presented in `monospaced font`.

Variable or non-literal text is presented in *italics*. For example: "Go to your *KVHOME* directory."

 **Note:**

Finally, notes of special interest are represented using a note block such as this.

1

Introduction to Oracle NoSQL Database

Welcome to Oracle NoSQL Database, providing multi-terabyte distributed storage for key-value pairs, with scalable throughput, and great performance. Oracle NoSQL Database services network requests to store and retrieve data, accessing data as either tables or key-value pairs. Oracle NoSQL Database services data requests with low latency, high throughput, and predictable data consistency, based on how you configure the store.

Oracle NoSQL Database uses Oracle Berkeley DB Java Edition as its underlying storage engine. For more information about Oracle Berkeley DB Java Edition, see [Oracle Berkeley DB Java Edition](#).

Oracle NoSQL Database offers full Create, Read, Update and Delete (CRUD) operations with adjustable durability guarantees. Oracle NoSQL Database is designed with high availability (HA), excellent throughput, and low latency, while requiring minimal administrative interaction.

Oracle NoSQL Database provides performance scalability. To increase performance, you can add hardware. If your performance is sufficient for your needs, you can purchase and manage fewer hardware resources.

Oracle NoSQL Database is designed for applications that require network-accessible data with user-definable read/write performance levels. A typical example is a web application servicing requests across the traditional three-tier architecture: web server, application server, and back-end database. In this configuration, Oracle NoSQL Database should be installed behind the application server, either taking the place of the back-end database, or working alongside it. To make use of Oracle NoSQL Database, you must supply code to run on the application server.

An application makes use of Oracle NoSQL Database by performing network requests against a data store, generically referred to as the KVStore. To make such data requests, you link the Oracle NoSQL Database driver to your application as a Java library (.jar file). Your code can then access any of the Java APIs that the library supplies. See *Java Direct Driver Developer's Guide*.

 **Note:**

Because Oracle NoSQL Database is tested using Java 8, use only that Java version with Oracle NoSQL Database.

You can also use a non-Java language driver to access table data stored in Oracle NoSQL Database. Drivers are available for C, C#, Node.js, and Python. For more information, see the specific Developer's Guide for the driver of interest.

Oracle NoSQL Database also provides SQL for Oracle NoSQL Database, which is an easy to use SQL-like language that supports read-only queries and data definition (DDL) statements. Use this SQL-like language to access table data for read-only queries and DDL statements.

To follow along with the query examples run with the interactive shell, see *SQL Beginner's Guide*.

To execute queries using the JAVA API, see *Java Direct Driver Developer's Guide*.

For a more detailed description of the SQL language (DDL, DML, and queries), see *SQL Reference Guide*.

NoSQL Database Server Licensing

Oracle NoSQL Database Server is available with two licensing options: Oracle NoSQL Database Community Edition (CE) and Oracle NoSQL Database Enterprise Edition (EE).

For a description on these two licenses, see [NoSQL Database Option Differences](#).

NoSQL Database Client Licensing

Oracle NoSQL Database client APIs are released as open source. Clients ship with source code and are released under the Apache 2.0 License. You use the client APIs to access Oracle NoSQL Database servers using either the Community Edition (CE) or Enterprise Edition (EE) licenses.

NoSQL Database Option Differences

Oracle NoSQL Database Server is available in two different options: Community Edition (CE), and Enterprise Edition (EE).

Community Edition (CE)

Community Edition is available using an open source license, and ships with source code.

Note:

The Community Edition is not always at the same release number as EE. For example, EE can be shipping version 19.1, while CE is available as 12.x.x version.

Enterprise Edition (EE)

Enterprise Edition requires a commercial license. It does not ship with source code.

Feature Differences between EE and CE

Oracle NoSQL Database Server Enterprise Edition includes new and updated features with each release. However, Community Edition is neither released as frequently as the Enterprise Edition, nor can it support every EE feature. Currently, CE does not support these features:

- Oracle GeoJSON Data support
- Kerberos Authentication Service integration

- Oracle Database External Table integration
- Oracle Coherence integration
- Oracle Event Processing integration - Streams Processor Engine
- Oracle Enterprise Manager integration
- Oracle Big Data Spatial and Graph integration
- Oracle Wallet integration for external password storage
- Oracle Semantic Graph integration

Architecture

Oracle NoSQL Database applications read and write data by performing network requests against an Oracle NoSQL Database data store, referred to as the KVStore. The KVStore is a collection of Storage Nodes, each of which hosts one or more Replication Nodes. Data is automatically spread across these Replication Nodes by internal KVStore mechanisms. Given a traditional three-tier web architecture, the KVStore either takes the place of your back-end database, or runs alongside it.

Optionally, a KVStore installation can be spread across multiple physical locations, each of which is called a *zone*. Zones are described in [Zones](#).

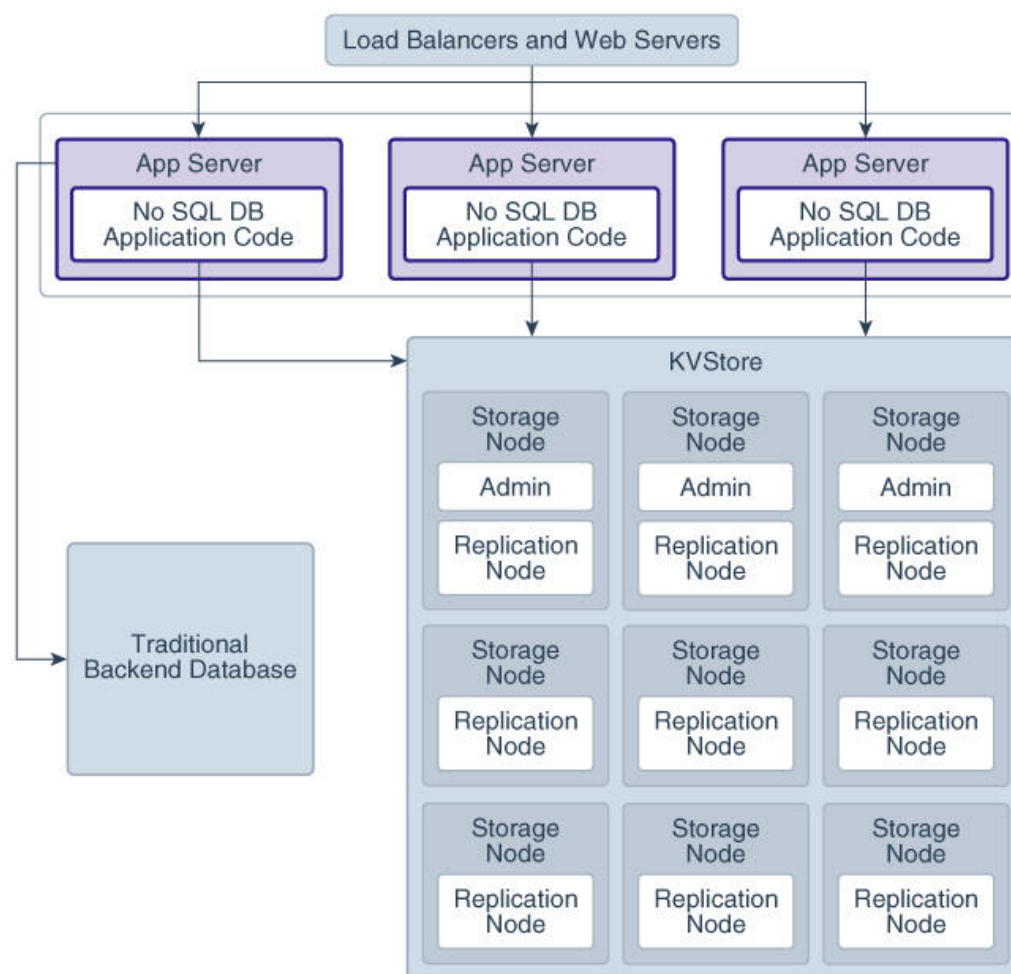
Note:

Replication Nodes are implemented using Berkeley DB, Java Edition (JE). JE is an enterprise-class, transaction-protected database, which is fully described in the Oracle Berkeley DB Java Edition.

The store contains multiple Storage Nodes. A *Storage Node* is a physical (or virtual) machine with its own local storage. The machine is intended to be commodity hardware. While not a requirement, each storage node is typically identical to all other Storage Nodes within the store.

The following illustration depicts a typical architecture used by an application that uses an Oracle NoSQL Database. Specifically, three of nine Storage Nodes each host an Admin process and a Replication Node. The remaining Storage Nodes each host a Replication node.

Figure 1-1 Typical Architecture for Oracle NoSQL Database Store



Every Storage Node hosts one or more Replication Nodes as determined by its *capacity*. A Storage Node's capacity serves as a rough measure of the hardware resources associated with it. Stores can contain Storage Nodes with different capacities, and Oracle NoSQL Database ensures that a Storage Node is assigned a proportional load size to its capacity.

A Replication Node, in turn, contains a subset of the store's data. Storage node data is automatically divided evenly into logical collections called *partitions*. Every Replication Node contains at least one, and typically many, partitions. Partitions are described in greater detail in [Partitions](#).

Finally, each Storage Node contains monitoring software that captures information ensuring the Replication Nodes that it hosts are running and healthy.

For more information on how to associate capacity with a Storage Node and know the best way to balance the number of Storage Nodes and Replication Nodes, see *Determining Your Store's Configuration* in the *Administrator's Guide*.

Replication Nodes and Shards

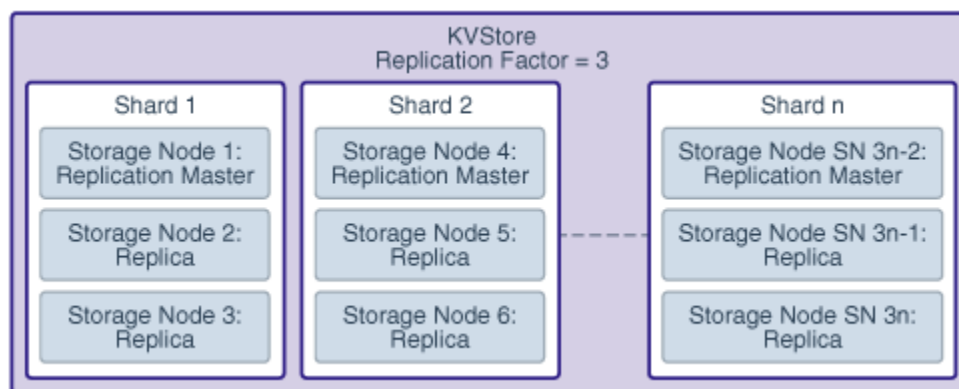
At a high level, you can think of a *Replication Node* as a single database containing tables or key-value pairs. Storage Nodes host one or more Replication Nodes. Because hosting a Replication Node depends on a healthy amount of resources, generally, Storage Nodes host only a single Replication Node. However, for installations with hardware that has abundant resources (memory, CPUs, and disks), Storage Nodes can, and do, host multiple Replication Nodes.

Your store's Replication Nodes are organized into *shards*. A single shard contains multiple Replication Nodes. Each shard has a *master node*. The master node performs all database write activities. Each shard also contains one or more read-only *replicas*. The master node copies all new write activity data to the replicas. The replicas are then used to service read-only operations.

While there can be only one master node per shard at any given time, any of the other shard members can become a master node. An exception to this is for nodes in a secondary zone as described below.

The following illustration shows how the KVStore is divided up into shards:

Figure 1-2 KVStore Shards



If the machine hosting the master node fails in any way, the master automatically fails over to one of the other nodes in the shard. One of the replica nodes is promoted automatically to master.

Production KVStores should contain multiple shards. At installation time you provide information that allows Oracle NoSQL Database to automatically decide how many shards the store should contain. The more shards that your store contains, the better your write performance is because the store contains more nodes that are responsible for servicing write requests.

Replication Factor

The number of nodes belonging to a shard is called its *Replication Factor*. The larger a shard's Replication Factor, the faster its read throughput, because there are more machines servicing the read requests. However, a large replication factor reduces write performance, because there are more machines to which writes must be copied.

A store can be installed across multiple physical locations called zones. You set a Replication Factor on a per-zone basis. Once you set the Replication Factor for each zone in the store, Oracle NoSQL Database makes sure the appropriate number of Replication Nodes are created for each shard residing in every zone in your store. Here are the terms used to describe these aspects of Oracle NoSQL Database:

- *Replication Factor*: The number of nodes belonging to a shard.
- *Zone Replication Factor*: The number of copies, or replicas, maintained in a zone.
- *Primary Replication Factor*: The total number of replicas in all primary zones.
- *Secondary Replication factor*: The total number in replicas in all secondary zones
- *Store Replication Factor* The total number of replicas in all zones across the entire store.

For additional information on how to identify the *Primary Replication Factor* and the implications of its value, as well on multiple zones and replication factors, see Replication Factor in the *Administrator's Guide*.

Partitions

All data in the store is accessed by one or more keys. A key might be a column in a table, or it might be the key portion of a key/value pair.

Keys are placed in logical containers called *partitions*, and each shard contains one or more partitions. Once a key is placed in a partition, it cannot be moved to a different partition. Oracle NoSQL Database distributes records evenly across all available partitions by hashing each record's key.

As part of your planning activities, you must decide how many partitions your store should have. You cannot configure the number of partitions after the store has been installed. For information about how to plan your store, see Initial Capacity Planning in the *Administrator's Guide*.

You can expand and change the number of Storage Nodes in use by the store. The store is then reconfigured to take advantage of the new resources by adding new shards. When this happens, existing data is spread across new and old shards by redistributing partitions from one shard to another. For this reason, it is desirable to have a large number of partitions to support fine-grained reconfiguration of your store.

As a general guideline, each shard should have at least 10 to 20 partitions. The number of partitions should be evenly divisible by the number of shards. Since the number of partitions cannot be changed after the initial deployment, plan the number of partitions for the maximum size of your store in the future. For example, while there is overhead in configuring a large number of shards, it is reasonable to specify a partition number that is 100 times the maximum number of shards you expect your store to contain.

Zones

A zone is a physical location that supports high capacity network connectivity between the Storage Nodes deployed within it. Each zone has some level of physical separation from other zones. Typically, each zone includes redundant or backup power supplies, redundant data communications connections, environmental controls (for example: air conditioning, fire suppression), and security devices. A zone can

represent a physical data center building, the floor of a building, a room, pod, or rack, depending on the particular deployment.

Oracle recommends installing and configuring your store across multiple zones. Having multiple zones provides fault isolation, and increases data availability in the event of a single zone failure. Multiple zones help mitigate systemic failures that affect an entire physical location, such as a large scale power or network outage.

There are two types of zones — *primary* and *secondary*. Primary zones are the default. They contain nodes that can serve as masters or replicas. Secondary zones contain nodes that can serve only as replicas. You can use secondary zones to make a copy of the data available at a distant location, or to maintain an extra copy of the data to increase redundancy or read capacity.

Only primary zones can have a Replication Factor equal to zero. Zero capacity Storage Nodes are used for Arbiter Nodes, which only primary zones can host.

You can use the command line interface to create and deploy one or more zones. Each zone hosts the deployed storage nodes. For additional information on zones and how to create them, see *Create a Zone* in the *Administrator's Guide*.

Arbiter Nodes

An `Arbiter Node` is a lightweight process that is capable of supporting write availability in two situations. First, when the primary replication factor is two and a single Replication Node becomes unavailable. Second, when two Replication Nodes are unable to communicate to determine which one of them is the master. The role of an Arbiter Node is to participate in elections and respond to acknowledge requests in these situations.

An Arbiter Node does not host any data. You create a Storage Nodes with zero storage capacity to host an Arbiter Node. While you can allocate Arbiter Nodes on Storage Nodes with a capacity greater than zero, those Arbiter Nodes have a lower priority during allocation than those on zero capacity Storage Nodes.

The Arbiter Node is allocated on a Storage Node outside of the shard. An error occurs if there are not enough Storage Nodes to host an Arbiter Node located on a different Storage Node from other shard members. The Arbiter Node provides write availability in the absence of a single Storage Node. The pool of Storage Nodes in a primary zone configured to host Arbiter Nodes is used for allocating an Arbiter Node.

For more information on Arbiter Nodes, see *Deploying an Arbiter Node Enabled Topology* in the *Administrator's Guide*.

Topologies

A *topology* is the collection of zones, storage nodes, shards, replication nodes, and administrative services that make up your NoSQL Database store. A deployed store has one topology that describes its state at a given time.

After initial deployment, the topology is laid out so as to minimize the possibility of a single point of failure for any given shard. This means that while a Storage Node might host more than one Replication Node, those Replication Nodes are never from the same shard. This improves the chances that the shard will have continuous availability for reads and writes, even if a hardware failure takes down the host machine.

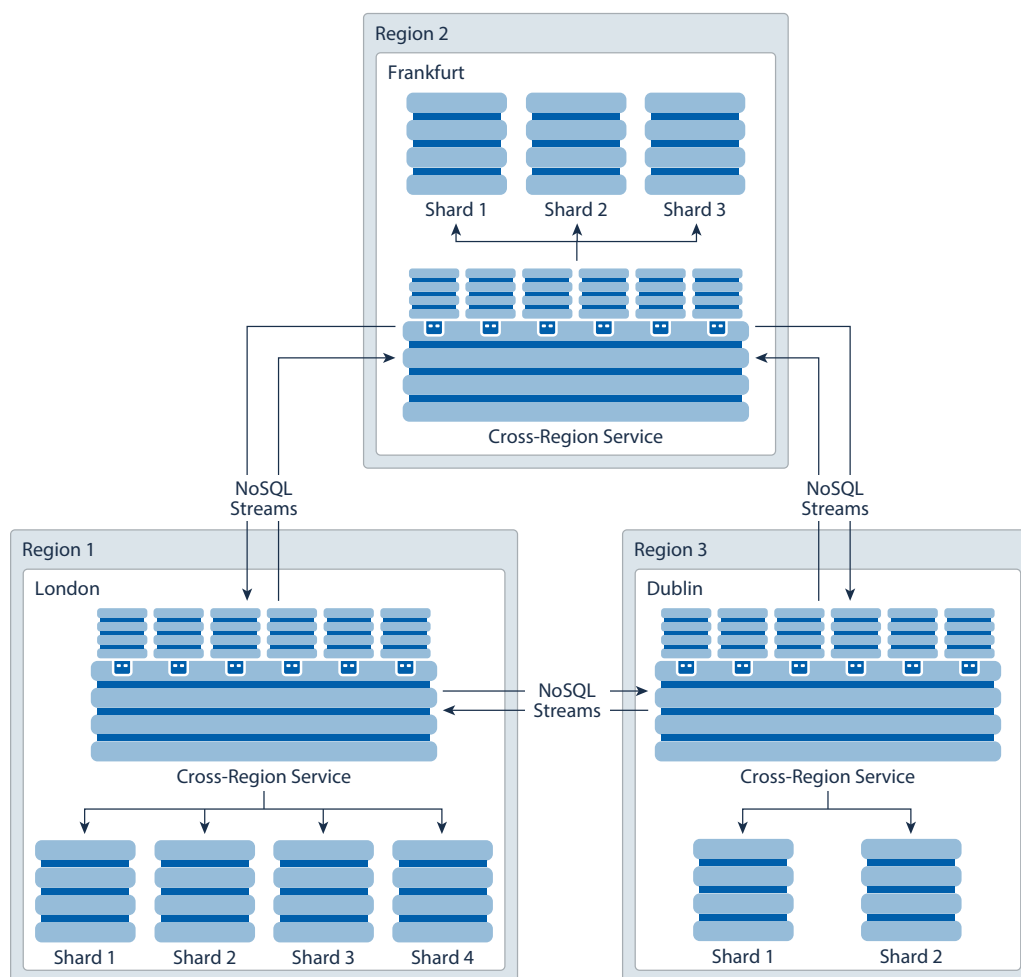
Arbiter Nodes are automatically configured in a topology if the primary replication factor is two and a zone is configured to host `Arbiter Nodes`.

Topologies can be changed to achieve different performance characteristics, or in reaction to changes in the number or characteristics of the Storage Nodes. Changing and deploying a topology is an iterative process. For information on how to use the command line interface to create, transform, view, validate and preview a topology, see *Steps for Changing the Store's Topology* in the *Administrator's Guide*.

Multi-Region Architecture

Oracle NoSQL Database applications read and write data by performing network requests against an Oracle NoSQL Database data store, referred to as the KVStore. Sometimes, organizations may need to set up multiple KVStores to maintain their NoSQL data. In more realistic situations, these KVStore clusters may even be geographically distributed. Oracle NoSQL Database multi-region architecture enables you to create tables in multiple KVStore clusters and maintain consistent data across these clusters.

For example, consider a use-case where an organization deploys three on-premise KVStores, one each at Frankfurt, London, and Dublin. In such a setup involving multiple KVStores, each independent Oracle NoSQL Database installation is referred to as a *Region*. Such an architecture having two or more independent, geographically distributed KVStore clusters bridged by bi-directional NoSQL Streams is known as *Multi-Region Architecture*.

Figure 1-3 Multi-Region Architecture

Suppose you want to collect and maintain similar data across multiple regions. You need a mechanism to create tables that can span across multiple regions and keep themselves updated with the inputs from all the participating regions. You can achieve these using Multi-Region tables. A *Multi-Region Table* or *MR Table* is a global logical table that is stored and maintained in different regions or installations. It is a *read-anywhere* and *write-anywhere* table that lives in multiple regions.

As you can see in the diagram, all the Multi-Region Tables defined in these regions are synchronized via NoSQL Streams. Essentially, all the distributed KVStore clusters form a fully connected graph. For each distributed KVStore cluster, there is one inbound stream from each remote KVStore cluster. This inbound stream subscribes the local KVStore to all the Multi-Region Tables from the remote KVStore. Each region must be running a *Cross-Region Service* (*XRegion Service*) to receive the data from the subscribed tables in the remote regions.

In addition, a list of points worth noting regarding the Multi-Region Architecture are:

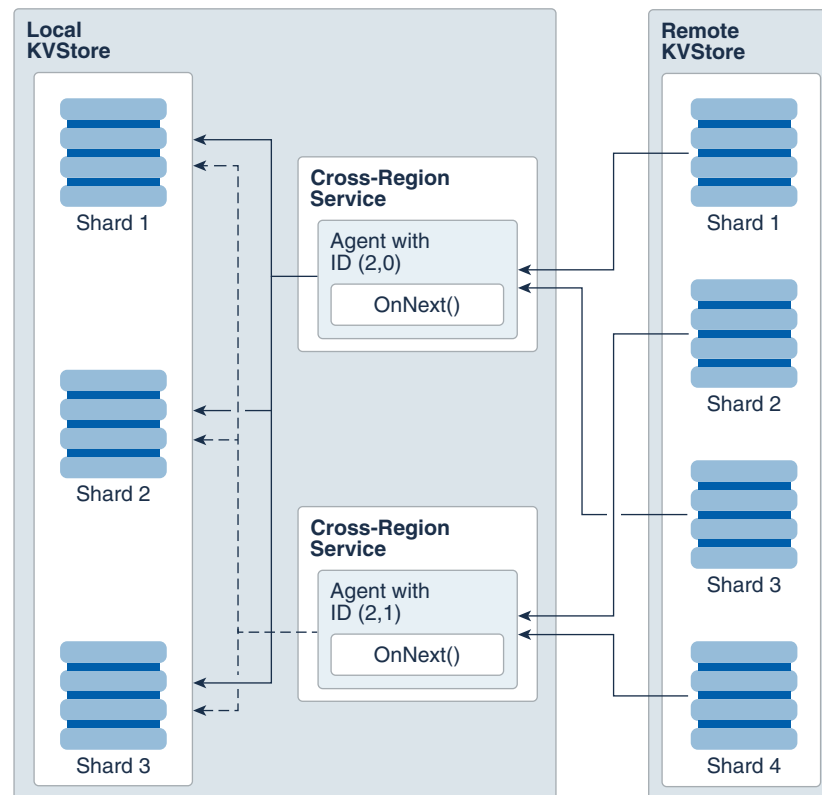
- The local and remote KVStores:
 - May have different topology
 - May experience elastic operations.

- Are independently managed, that is, each KVStore has its own index, security credentials etc.
- Have sufficient create and read table privileges to each other.
- The inbound and outbound streams are:
 - Completely symmetrical.
 - Independently managed without any coordination between the outbound and the inbound streams.

Cross Region Service

In a Multi-Region Oracle NoSQL Database setup, a *Cross-Region Service* or *XRegion Service* is a standalone service running on a separate node. In simple terms, this is also called an *agent*. The XRegion Service is deployed when you are connecting the local KVStore with a remote KVStore to create a Multi-Region Table.

Figure 1-4 Inbound Stream from Remote KVStore



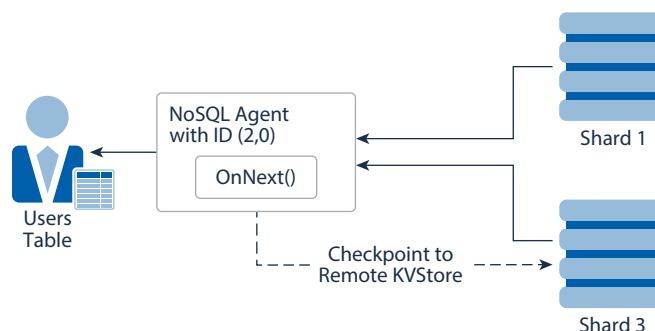
Consider an example where the remote KVStore has four shards and the local KVStore has three shards. The local KVStore deploys two NoSQL agents to stream two shards each from the remote KVStore, as depicted in the diagram above. These agents subscribe to updates from the remote KVStore and publish the new and modified rows to the local KVStore. All the agents connecting a remote KVStore to a local KVStore are referred as the *Agent Group* for the local KVStore.

Note:

Eventhough the above example demonstrates *Agent Groups*, please note that the current release of Oracle NoSQL Database does not support configuring multiple agents. As of the current release, you can configure only one agent for each local region.

As you can see in the diagram below, an agent enables streaming the data from a remote KVStore to a set of MR Tables in the local KVStore.

Figure 1-5 View of a Single Agent



Each inbound stream utilizes the subscriber group feature in the Streams API to create a group of subscribers to stream from a store. Each local agent is responsible for:

- Establishing the inbound subscription stream from a remote KVStore.
- Maintaining the connection and reconnect during any failures.
- Checkpointing the subscription stream in case of any failures.
- Subscribing writes to the MR Tables from a remote KVStore.
- Automatically dealing with elastic operations in the remote KVStore.

For a local KVStore, the overhead of inbound stream consists of:

- A group of threads in the NoSQL agent that streams the data from the remote KVStore. Please note that these threads run outside the local store as a standalone agent. Even though the agent serves the local KVStore, it does not add to the local KVStore's expense.
- Local PUT or DELETE resulting from the streamed data from the remote KVStore after conflict resolution.

For a local KVStore, the overhead of outbound stream involves the create, read, and write checkpoints for the remote KVStore.

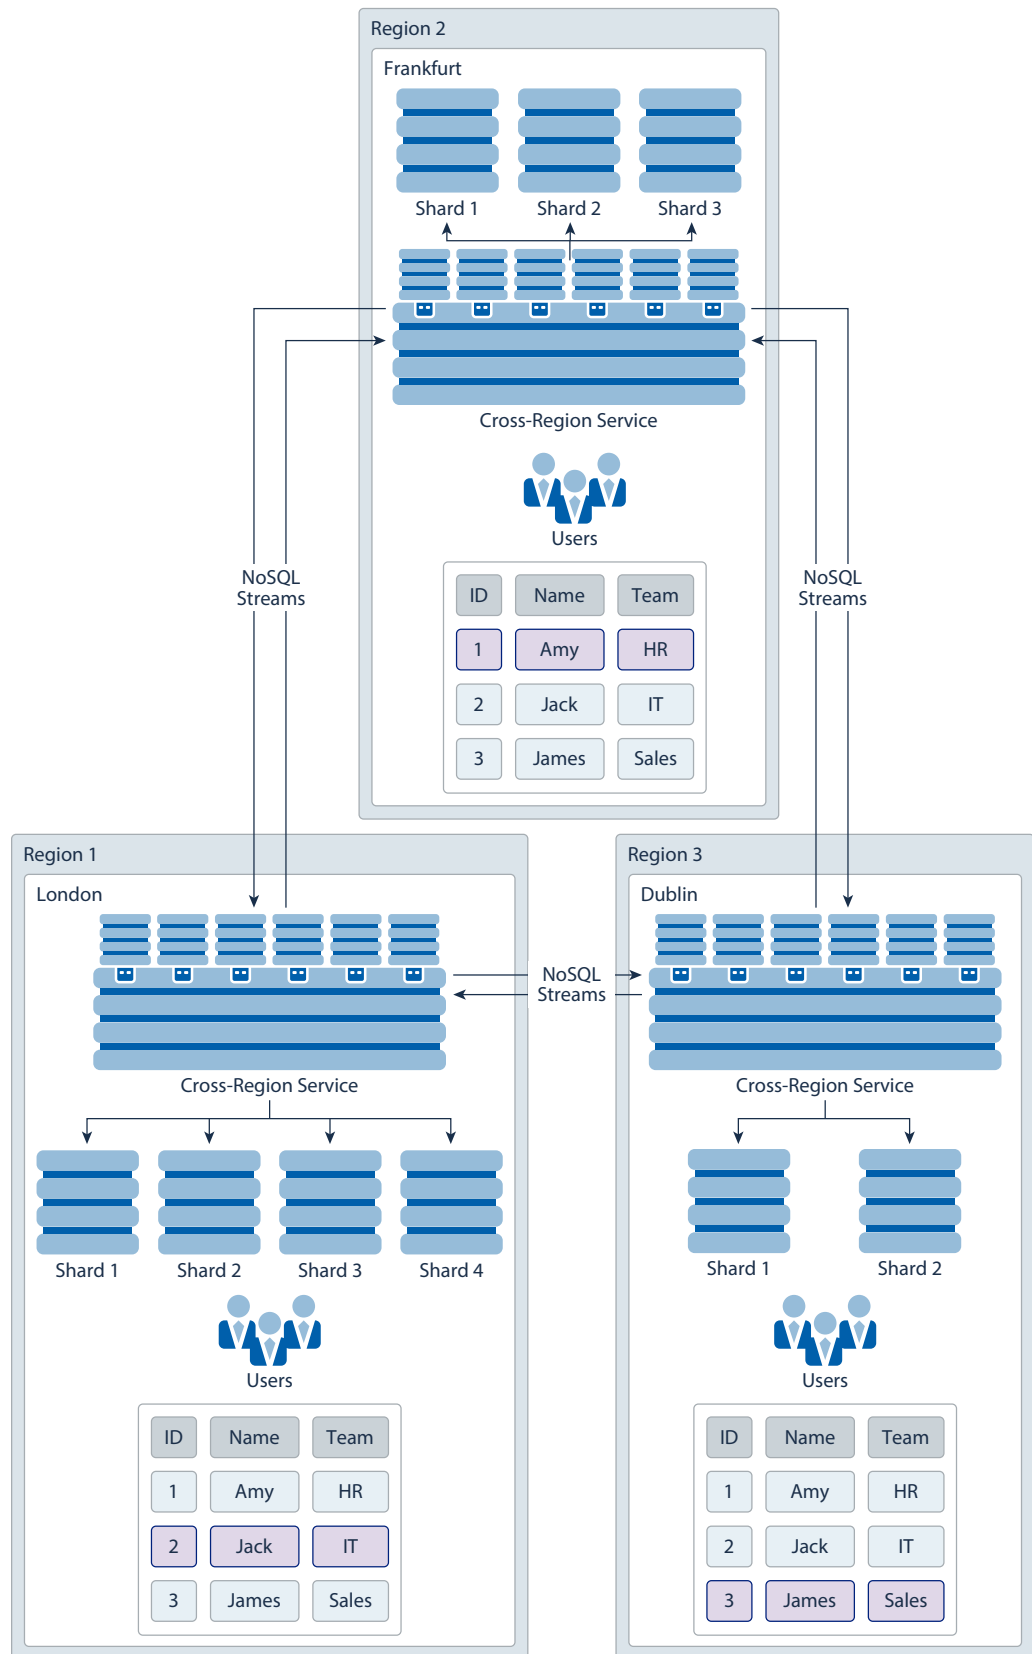
Life Cycle of Multi-Region Tables

To create and use Multi-Region Tables (MR Tables) in Oracle NoSQL Database, you must be aware of the sequence of tasks to execute and the related concepts.

For clarity, let us discuss the life cycle of MR Tables with an example. Consider an Oracle NoSQL Database with three regions, Frankfurt, London, and Dublin. Assume

that you want to create a table called `Users` to store the user details for all the three regions as depicted in the diagram.

Figure 1-6 Multi-Region Table



The sequence of tasks that you must perform to create and manage the `Users` table as a Multi-Region Table are:

- **Deploy Independent KVStores:** You must deploy KVStore in each region of the Multi-Region NoSQL Database independently. See Configuration Overview in the *Administrator's Guide*.
- **Set Local Region Name:** After deploying the KVStore and before creating the first MR Table in each participating region, you must set a local region name. You can change the local region name as long as no MR Tables are created in that region. After creating the first MR Table, the local region name becomes immutable. The local region name is completely independent of the KVStores created in that region. See Set Local Region Name in the *Administrator's Guide*.
- **Configure XRegion Service:** Before creating any MR Table, you must deploy an XRegion Service with one or more agents. The agent runs independently with the local KVStore and it is recommended to deploy it close to the local KVStore. While setting up a secure KVStore to support multi-region tables, the administrator has to grant the following permissions to the XRegion Service agent:
 - `WRITE_SYSTEM_TABLE` (or the equivalent `writesystable` role) to the local store.
 - Write permission on all the multi-region tables in the local store.
 - Read permission on all the multi-region tables in the remote stores.
 - Write permission for checkpoint table in the remote stores.

To learn how to deploy an agent, see Configure XRegion Service in the *Administrator's Guide*.

- **Start XRegion Service:** You must start XRegion service in each region using the `XRSTART` command. As this service is a long-running process, it is recommended to invoke it as a background process by appending the `&` at the end of the command. see Start XRegion Service in the *Administrator's Guide*.

 **Note:**

The local KVStore must be started before starting the XRegion Service. If the KVStore in the local region has not started or is not reachable, the XRegion Service will not start.

- **Create Remote Regions:** Before creating and operating on the MR table, you must define the remote regions. A remote region signifies a region different from the region where the command is executed. In this example, to create the MR Table called `Users` from the Frankfurt region, you must first define the other two regions, that is, London and Dublin using the `CREATE REGION` DDL command. To learn how to create remote regions, see Create Remote Regions in the *Administrator's Guide*.
- **Create the MR Table:** You must create an MR Table on each KVStore in the connected graph, and specify the list of regions that the table should span. In this example, to create the `Users` table as an MR Table at the Frankfurt and London regions, you need to execute the `CREATE TABLE` command specifying Frankfurt and London as the regions. The order in which you list the regions in the DDL command does not matter. After you create the `Users` MR Table, it will be included in the incoming stream from each remote KVStore specified. Symmetrically at the

remote KVStore the `Users` table will be included in its own incoming stream too. To create the MR table successfully, you must:

- Ensure that you acquire the necessary privileges to create the table in the specified regions, in advance. Otherwise, the MR Table creation will fail in all the regions. See KVStore Required Privileges in the *Security Guide*.
- Specify at least one region in the `CREATE TABLE` DDL command. If you specify only one region, then the MR Table is created only in the specified region and no writes will be replicated to the other regions.

 **Note:**

Even though a single-region MR table works similar to a local table, the difference between them is that the single-region MR Table can be expanded to multiple regions in future.

To learn how to create an MR Table, see Create MR Tables in the *Administrator's Guide*.

- **Perform Read/Write Operations on the MR Table (Optional):** After creating the MR Table, you can perform read or write operations on the table using the existing data access APIs or DML statements. There is no change to any existing data access APIs or DML statements to work with the MR Tables. The following aspects applicable to the regular tables apply to the MR Tables also without any deviation:
 - **Durability and consistency configurations and constraints:** For any local writes to an MR table, the semantics of consistency model does not change. It is the same as any writes to a regular (non MR) table.

 **Note:**

In case of MR Tables, absolute consistency is not global across the participating regions. It is only local to a single region where you perform the read and write operations.

- **Table index infrastructure:** Creating primary and secondary indices on the MR Table in each region remains the same as with any regular (non MR) table. However, if you wish to drop an MR Table from any region, you must first drop all the indices defined on this table.

Read Operations:

- Each Read operation on an MR Table is a local read, that is, you read only the local copy of the data. However, this local copy may have rows that might have come from one of the other participating regions, as a result of a table sync-up via Oracle NoSQL Streams.

Write Operations:

- Whenever you execute a write operation (`INSERT`, `UPDATE`, or `DELETE`) on an MR Table, the changes will be replicated across multiple regions asynchronously. It means, when you write a row in the local region, the write operation is executed completely in the local region without waiting for the subscribing regions to update.

- The latency for replicating the changes across multiple regions includes the time taken to:
 - * Complete the write operations at the remote region, and
 - * Receive the data from the subscribed tables.
- If multiple regions update a row with the same primary key, a built-in conflict resolution rule is applied to decide which region's update is considered as final. In all such cases, this built-in conflict resolution rule will cause the update with the latest timestamp to win and commit to the database.
- The above mentioned built-in conflict resolution rule applies to the TTL value updates too.
- When you delete a row from an MR Table, the system allows time to ensure that the change is propagated to all the other regions where this table exists.
- You can define a TTL value while inserting or updating a row in the MR Table. This value applies only to the row being added or updated. The row-level TTL overrides table level TTL if any exists.
- An MR Table can have different table level TTL values in different regions.
- When a row is replicated to other regions, its expiration time is replicated as an absolute timestamp to the replicated rows. This can be either the default table level TTL value or a row level override that is set by your application. Therefore, this row will expire in all the regions at the same time, irrespective of when they were replicated.
- If a row expires before it makes it to one of the regions during replication, the rows that are already replicated to other regions will expire immediately after persistence.

See Access and Manipulate MR Tables in the *Administrator's Guide* for examples.

- **Add New Regions to the MR Table (Optional):** Oracle NoSQL Database lets you expand an MR Table to new regions. It effectively means adding new regions to an existing MR Table. In the example being discussed, suppose you created the `Users` table only in two regions, Frankfurt and London. Later, if you want to expand this `Users` table to the Dublin region, you must:
 - Create the `Users` MR Table in the new region, that is, Dublin. Note that you must specify all the three regions while creating the MR Table in the new region. See Create MR Table in New Region in the *Administrator's Guide*.
 - Add the new region (Dublin) to the `Users` MR Table in existing regions, that is, Frankfurt and London. This is achieved with the help of the `ALTER TABLE` DDL command. See Add New Region to Existing Regions in the *Administrator's Guide*.

 **Note:**

Depending on the volume of the data in the existing regions, it might take some time to initialize the MR Table in the new region with the data from the other regions. However, the MR Table in the new region is available for read/write operations immediately after its creation.

To learn how to expand an MR Table with detailed code demonstrations, see Use Case 2: Expand a Multi-Region Table in the *Administrator's Guide*.

- **Remove an Existing Region from the MR Table (Optional):** Not only can you add new regions to an existing MR Table but also remove any regions linked to it. It effectively means that you disconnect the MR Table from a particular region so that MR Table is not synchronized with any writes from the removed region. This is called contracting an MR Table. To learn how to contract an MR Table with detailed code demonstrations, see Use Case 3: Contract a Multi-Region Table in the *Administrator's Guide*.
- **Drop Remote Regions (Optional):** You can drop one or more participating regions from a Multi-Region Oracle NoSQL Database setup as per your business requirement. However, before removing a region from a Multi-Region NoSQL Database, it is recommended to:
 - Stop writing to all the MR Tables linked to that region.
 - Ensure that all writes to the MR Tables in that region have synchronized with other regions. This helps in maintaining consistent data across the different regions.

**Note:**

Even though NoSQL Database lets you drop a region directly, it is a recommended practice to isolate that region from all the other regions before dropping it. This ensures that the existing regions are no longer linked with the region being dropped.

To learn how to drop a region from a Multi-Region NoSQL Database, see Use Case 4: Drop a Region in the *Administrator's Guide*.

- **Shut Down XRegion Service and KVstores:** In a case where you want to relocate your XRegion Service to another host machine, you must shut it down in the current machine and then restart it in the new host machine. See Stop XRegion Service in the *Administrator's Guide*.

Data Models

You can model your data in Oracle NoSQL Database by using Tables or a key-value interface.

Tables are the easiest way to model data. They provide the highest level of abstraction, they are simple to model and should be familiar to any developer. This model also supports secondary indices and table evolution.

If you are using tables, then you can use JSON to model data. If strongly typed data is not a priority, then this is a good choice.


Oracle NoSQL Database also supports multi-region tables. A multi-region table is a global logical table that is stored and maintained in different regions or installations. In short, they are called as MR Tables.

Finally, if you want to serialize data, manage the key structure, manage secondary indices through index views, manage evolution and security through your client code, or work with large objects, then you can use the key-value interface.

Consistency

Oracle NoSQL Database provides several different consistency policies. At one end of the spectrum, applications can specify absolute consistency, which guarantees that all reads return the most recently written value for a designated key. At the other end of the spectrum, applications capable of tolerating inconsistent data can specify weak consistency, allowing the database to return a value efficiently even if it is not entirely up to date. In between these two extremes, applications can specify time-based consistency to constrain how old a record might be or version-based consistency to support both atomicity for read-modify-write operations and reads that are at least as recent as the specified version.

The following illustration depicts the range of consistency policies that can be used by an application that makes use of Oracle NoSQL Database:



Flexible consistency policies enables developers to easily create business solutions providing data guarantees while meeting application latency and scalability requirements.

Durability

Oracle NoSQL Database provides a range of durability policies that specify what guarantees the system makes after a crash. At one extreme, applications can request that write requests block until the record has been written to stable storage on all copies. This has obvious performance and availability implications, but ensures that if the application successfully writes data, that data will persist and can be recovered even if all the copies become temporarily unavailable due to multiple simultaneous failures. At the other extreme, applications can request that write operations return as soon as the system has recorded the existence of the write, even if the data is not persistent anywhere. Such a policy provides the best write performance, but provides no durability guarantees.

The following illustration depicts the range of durability policies that can be used by an application that makes use of Oracle NoSQL Database:



By specifying when the database writes records to disk and what fraction of the copies of the record must be persistent (none, all, or a simple majority), applications can enforce a wide range of durability policies.

Quorum

Operations that modify data in Oracle NoSQL Database require that at least a simple majority of primary nodes be available to form a *quorum* in the shard that stores the specified key.

Quorum is the minimum number of primary nodes required in a shard, or in the set of admin nodes, to permit electing a master to support write operations. To form a quorum requires that a minimum number of primary nodes represent a majority in the group.

 **Note:**

Secondary nodes are not counted when computing the quorum.

Consider the following example using a store with four zones. Zones 1, 2, and 3 are primary zones with replication factor 1, and zone 4 is a secondary zone with replication factor 1. The number of primary nodes in each shard is 3, which is the sum of the replication factors for the primary zones. In a group of 3 nodes, 2 is the smallest number of nodes that represent a majority, so the quorum is 2. The secondary nodes in zone 4 have no impact on the quorum.

In general, to compute the quorum, first determine the primary replication factor, which is the sum of the replication factors of all primary zones. The quorum value must be one greater than half of the primary replication factor, rounding down when computing the half.

For example, for primary replication factor of 1, the quorum is 1. For primary replication factor of 5 the quorum is 3. For primary replication factor of 6, the quorum is 4.

Administration

The Administration command line interface (CLI) is the primary tool you use to manage your store. You use it to configure, deploy, and change store components. Use the CLI to verify the system, check service status, check for critical events and browse the store-wide log file.

You can use the CLI to get, put, and delete store records or tables, retrieve schema, and display general information about the store. It can also be used to diagnose problems or potential problems in the system, fix actual problems by adding, removing or modifying store data and/or verify that the store has data. The CLI is particularly well-suited for a developer who is creating an Oracle NoSQL Database application. Developers can use the CLI to populate a store with a small number of records to use for development purposes, or to examine the store's state as part of debugging activities.

Access the command line interface using this command:

```
java -Xmx64m -Xms64m -jar KVHOME/lib/kvstore.jar runadmin
```

 **Note:**

To avoid using too much heap space, specify `-Xmx` and `-Xms` flags for Java when running administrative and utility commands.

For a complete list of all CLI commands and their usage, see Admin CLI Reference in the *Administrator's Guide*.

KVLite

KVLite is a simplified version of Oracle NoSQL Database. It provides a single-node store that is not replicated. It runs in a single process without requiring any administrative interface. You configure, start, and stop KVLite using a command line interface.

KVLite is intended for use by application developers who need to unit test their Oracle NoSQL Database application. It is not intended for production deployment, or for performance measurements.

KVLite is installed when you install KVStore. It is available in the `kvstore.jar` file in the `lib` directory of your Oracle NoSQL Database distribution.

For more information on KVLite, see *Quick Start to KVLite*.

Saving Admin CLI History

By default, Oracle NoSQL Database uses the Java Jline library to support CLI history that you can save. To disable this feature, set the following Java property while starting the `runadmin` program:

```
java -Xmx64m -Xms64m \  
-Doracle.kv.shell.jline.disable=true -jar KVHOME/kvstore.jar \  
runadmin -host <hostname> -port <portname>
```

Unless you disable the feature, CLI history is saved to a file so that it is available after restart. By default, Oracle NoSQL Database attempts to save 500 lines of history in the following file, which is created and opened automatically:

```
KVHOME/.jlineoracle.kv.impl.admin.client.CommandShell.history
```

 **Note:**

If the Admin CLI cannot open the history file, it fails silently. The CLI runs without saving any history.

To change the default history file path, set the location of the `oracle.kv.shell.history.file="path"` Java property.

To change the default number of lines, set the value of the `oracle.kv.shell.history.size=<int_value>` Java property.

Monitoring

Information about the performance and availability of your store is available. You can monitor the information through an API class, log files, and Java Management Extensions (JMX).

These agents provide interfaces on each Storage Node that allow management clients to poll them for information about the status, performance metrics, and operational parameters of the Storage Node and its managed services, including replication nodes, and admin instances. Also, JMX can be used to monitor Arbiter Nodes.

For more information, see Java Management Extensions (JMX) in the *Administrator's Guide*.

Troubleshooting

Errors can occur in your store deployment. Tools, commands, logs and procedures can be used in order to solve problems.

To catch configuration errors early, you can use the `Diagnostics Utility`. You can also use this tool to package information and files to send them to Oracle Support, for example.

For more information on troubleshooting your store, see Troubleshooting in the *Administrator's Guide*.

Access and Security

There are two ways to access the KVStore and its data.

For routine access to the data, use Java APIs that application developers use to allow applications to interact with the Oracle NoSQL Database Driver. The driver communicates with the store's Storage Nodes to perform whatever data access the developer application requires.

For administrative access to the store, use the command line interface (CLI). System administrators use this interface to perform any actions that are required by Oracle NoSQL Database. You can also monitor the store using the CLI interface.

For most production stores, authentication over SSL is normally required by both the command line interface and the Java APIs. While you can install a store that does not require authentication, this is not recommended. For details on Oracle NoSQL Database's security features, see the *Security Guide*.

 **Note:**

Oracle NoSQL Database is intended to be installed in a secure location where physical and network access to the store is restricted to trusted users. For this reason, at this time Oracle NoSQL Database's security model is designed to prevent accidental access to the data. It is *not* designed to prevent denial-of-service attacks.

Integration

Oracle NoSQL Database can be integrated with Apache Hadoop and products in the Oracle stack. The following sections describe more about integration.

Oracle Database Mobile Server Integration

As of Oracle Database Mobile Server (Oracle DMS) 12.1 release, Oracle NoSQL Database can be integrated with Oracle Database Mobile Server. Oracle DMS facilitates the development, deployment and management of mobile database applications for a large number of mobile users.

The main integration feature is synchronizing Oracle NoSQL with Oracle Berkeley DB/SQLite/Java DB.

Oracle DMS can be used as a data synchronization engine between Oracle NoSQL and mobile client databases including Oracle Berkeley DB, SQLite and Java DB. Existing tables in an Oracle NoSQL database can be published to Oracle DMS using publish/subscribe APIs. Oracle DMS creates corresponding tables on the client-side and allows them to synchronously upload their data to Oracle NoSQL. Since, Oracle DMS completely manages the data format conversion between the local client database and Oracle NoSQL DB, no additional coding or scripting is required. The mobile to NoSQL paradigm fits best where the number of clients is large, and/or the amount of collected data is large, and/or the server application requires NoSQL as data storage.

With Oracle NoSQL as data storage, mobile and embedded applications can use Oracle DMS as a complete and reliable solution for synchronizing their data to the highly scalable Oracle NoSQL database. For more information, see *Customizing Synchronization With Your Own Queues* in the *Oracle Database Mobile Server Developer's Guide*. You can also find the FleetControl sample application in Oracle Mobile Development Kit (a development toolkit for Oracle DMS), which demonstrates how to use Oracle NoSQL as a back-end database and how to synchronize application data between mobile client database and Oracle NoSQL via Oracle DMS.

Hadoop Integration

Oracle NoSQL Database can be integrated with Apache Hadoop systems using the `oracle.kv.hadoop.KVInputFormat` class. This class allows you to read data from Oracle NoSQL Database and then prepare it for insertion into a Hadoop system. To move data back to your Oracle NoSQL Database, you can read data from the Hadoop system using the standard mechanisms, and then write the records to Oracle NoSQL Database using the APIs. See *Integration with Apache Hadoop MapReduce* in the *Integrations Guide*.

An example of using `KVInputFormat` to read data from Oracle NoSQL Database in a Map/Reduce job can be found in the `<KVHOME>/examples/hadoop` directory.

Property Graph Integration

Oracle Big Data Spatial and Graph can be configured to use Oracle NoSQL Database to support its property graph feature. This feature supports graph operations, indexing, queries, search and in-memory analytics.

Graphs are commonly used to model, store, and analyze relationships found in social networks, cyber security, utilities and telecommunications, life sciences and clinical data, and knowledge networks. Typical graph analyses encompass graph traversal, recommendations, finding communities and influencers, and pattern matching.

For more information, see *Using Property Graphs in a Big Data Environment in the Oracle Big Data Spatial and Graph User's Guide and Reference*.

Oracle External Tables Integration

Oracle NoSQL Database data can be accessed using Oracle Database's External Tables feature. This capability allows NoSQL Database data to be read into Oracle Database. Oracle NoSQL Database data cannot be modified using the External Tables feature.

Note that this is a feature which is only available to users of the Oracle NoSQL Database Enterprise Edition.

To use the Oracle Database External Table feature to read Oracle NoSQL Database data, you must use the `<KVHOME>/exttab/bin/nosql_stream` preprocessor to populate our Oracle tables with the data. You must then configure your Oracle Database to use the External Tables feature.

For information on how to use the `nosql_stream` preprocessor, and how to configure Oracle Database to use External Tables, see [oracle.kv.exttab package summary](#) in the *Java Direct Driver API Reference*.

Coherence Integration

Oracle Coherence is a middleware application that provides data management support for clustered applications. The data that an application delegates to Oracle Coherence are automatically available to and accessible by all servers in the application cluster. By distributing data across multiple machines, Oracle Coherence solves problems related to achieving availability, reliability, scalability, performance, serviceability and manageability of clustered applications.

For more information on Oracle Coherence, see the Oracle Coherence Documentation.

To provide these solutions, Oracle Coherence implements a cache. This cache can be customized to use a number of different data repositories to store the cached data. One such data repository is Oracle NoSQL Database.

Note that this is a feature which is only available to users of the Oracle NoSQL Database Enterprise Edition.

To integrate with Oracle NoSQL Database, Oracle Coherence must be customized using a combination of configuration XML, stock Coherence code, and custom Oracle

NoSQL Database code. The Oracle NoSQL Database code is implemented using classes provided by the `oracle.kv.coherence` package.

Oracle NoSQL Database can be used to support two different caching strategies for Oracle Coherence. The first of these is implemented by [oracle.kv.coherence.NoSQLBinaryStore](#). See the *Java Direct Driver API Reference*. This class allows you to implement cache data that is not meant to be shared with non-cache-based applications, and so uses a data format that is fairly opaque. This is an efficient and easy-to-configure caching option.