

Oracle NoSQL Database

Compared to MongoDB

Overview

- Oracle NoSQL Database and MongoDB server are both licensed under AGPL while MongoDB has certain client drivers under the Apache 2.0 license.
- Oracle NoSQL Database is in many respects, as a NoSQL Database implementation leveraging BerkeleyDB in its storage layer, a commercialization of the early NoSQL implementations which lead to the adoption of this category of technology. Several of the earliest NoSQL solutions were based on BerkeleyDB and some are still to this day e.g. LinkedIn's Voldemort. The Oracle NoSQL Database is a Java based key-value store implementation that supports a value abstraction layer currently implementing Binary and JSON types. Its key structure is designed in such a way as to facilitate large scale distribution and storage locality with range based search and retrieval. The implementation uniquely supports built in cluster load balancing and a full range of transaction semantics from ACID to relaxed eventually consistent. In addition, the technology is integrated with important open source technologies like Hadoop / MapReduce, an increasing number of Oracle software solutions and tools and can be found on Oracle Engineered Systems.
- MongoDB is an implementation of a key-value store that supports the single value abstraction JSON. The focus of MongoDB is on developer friendliness and it has a rich set of proprietary API's layered on top of a secondary index support for the JSON value type. MongoDB has an increasingly complete set of query capabilities approaching the kind of functionality found in traditional relational databases, using a declarative API rather than an expression language like SQL. MongoDB is known for being easy to get started with due to its integration with many scripting languages making database accessibility less cumbersome than compiled language integration. MongoDB is a shard based clustering database which very recently introduced consistent hashing to compliment its user defined range based implementation.

Comparison

The table below gives a high level comparison of Oracle NoSQL Database and MongoDB features/capabilities. Low level details are found in links to Oracle and MongoDB online documentation.

Feature/Capability	Oracle NoSQL Database	MongoDB
Data Model	<p>Oracle NoSQL Database has a flexible key-value data model leveraging a value abstraction layer. The value abstractions supported at this time are Binary and JSON(Avro). A table-structure value abstraction is coming soon</p> <ul style="list-style-type: none">• Record Design Considerations• Avro Schemas	<p>MongoDB's data format is BSON (binary equivalent to JSON) stored as documents (self-contained records with no intrinsic relationships). Documents in MongoDB may store any of the defined BSON types and are grouped in collections.</p> <ul style="list-style-type: none">• Documents• Data Types and Conventions
Storage Model	<p>Oracle NoSQL Database storage model is a write ahead logging implementation proven in millions of BerkeleyDB deployments. It's an append only implementation that enables efficient write throughput with background compaction for space reclamation. Write operation durability can be controlled by the user to allow multi-memory write operations without fsync or with fully durable disk sync. Data is partitioned into a fix space that has logical overlays. So, data in partitions can move between logical shard</p>	<p>MongoDB's default storage system is the Memory-Mapped Storage Engine. It uses memory mapped files for all disk I/O. It is the responsibility of the OS to manage flushing data to disk and paging data in and out.</p> <ul style="list-style-type: none">• Caching

	<p>representations, but must be moved at the granularity of these partitions.</p> <ul style="list-style-type: none"> • BDB Storage - NoSQL before NoSQL was cool • The evolution of BerkeleyDB 	
<p>Data Access and APIs</p>	<p>Oracle NoSQL Database has client library API's for Java and C. In the works are a Command Line Interface and Javascript API.</p> <ul style="list-style-type: none"> • Client APIs 	<p>MongoDB uses a custom, socket-based wire protocol with BSON as the interchange format.</p> <ul style="list-style-type: none"> • Mongo Wire Protocol <p>10Gen and the Mongo community support many client libraries.</p> <ul style="list-style-type: none"> • Client-Libraries
<p>Query Types and Query-ability</p>	<p>Oracle NoSQL Database provides key access methods (put, get, delete) including multi-key variations with large result set streaming support.</p> <p>The database can also be accessed using SQL as an external table from within a relational database.</p> <p>It is integrated with and can participate in MapReduce operations from a Hadoop environment.</p> <ul style="list-style-type: none"> • Searching in Oracle NoSQL • External Table Support • NoSQL and 	<p>MongoDB has a query interface that has some similarities to relational databases, including secondary indexes that can be derived from the stored documents. MongoDB also has a facilities for performing MapReduce queries and ad-hoc queries on documents. Hadoop support is available, too.</p> <ul style="list-style-type: none"> • Querying • Indexes • MapReduce • MongoDB Hadoop Adapter

	<p>MapReduce</p> <ul style="list-style-type: none"> • Using Range Queries 	
Data Versioning and Consistency	<p>Oracle NoSQL Database provides control at the operation level for consistency and durability. Each operation can be fully ACID, flushing and syncing all data to disk before taking quorum on the operation to allowing a fire and forget into local or remote memory. Read consistency is obtained thru quorum control spanning the range of requiring all holders of a copy of data to agree to just getting the result from a first responder. This provides the ultimate control for the developer of both transactional and eventually consistent applications.</p> <p>Flexible Consistency options</p>	<p>MongoDB exhibits strong consistency. Eventually consistent reads can be accomplished via secondaries. A MongoDB cluster (with auto-sharding and replication) has a master server at a given point in time for each shard.</p> <ul style="list-style-type: none"> • On Distributed Consistency
Concurrency	<p>Oracle NoSQL Database concurrency is controlled thru replication groups with an elected master. Reads can be serviced from any node in a replication group and writes are performed at the currently elected master, then replication chained to the replicas in the group. Read consistency is tied to concurrency, controlled by quorum, version, timestamp,</p>	<p>MongoDB relies on locks for consistency. As of version 2.2, MongoDB has a DB Level Lock for all operations.</p> <ul style="list-style-type: none"> • Locks • DB Level Locking • How Does Concurrency Work?

	<p>all.</p> <p>Durability Guarantees</p>	
Replication	<p>Oracle NoSQL Database supports replication for both availability and scalability. It uses a consistent hashing algorithm over a fixed, highly granular, partition definition. Partitions are replicated in groups according to latency demands of the application, configured by a replication factor.</p> <ul style="list-style-type: none"> • Replication configuration <p>There is a topology aware driver that is linked with the client application. Writes use the driver to hash inserts to the currently elected master and then a cascading replication occurs to the replicas belonging to the replication group where that master resides. How many data replications must occur and whether or not those replications are to memory space or disk for the respective replica can be configured on a per operation basis.</p> <ul style="list-style-type: none"> • Topologies 	<p>Mongo manages replication via replica sets, a form of asynchronous master/slave replication. Traditional master/slave replication is available but not recommended.</p> <ul style="list-style-type: none"> • Replication • Replica Sets • Master/Slave
Scaling Out and In	<p>Oracle NoSQL Database scales out by redistribution of data partitions to newly added</p>	<p>Mongo relies on sharding for scaling out. This involves designating a certain server to hold certain chunks of the</p>

	<p>hardware resources. When new hardware is added to the system, an administrator, via a browser based console or CLI, can issue a request to rebalance the cluster. The administrator has the option of just letting it go or throttling or running during certain windows of time, pausing the process, etc.</p> <ul style="list-style-type: none"> • Managing Topology Changes 	<p>data as the data set grows.</p> <ul style="list-style-type: none"> • Sharding in MongoDB • Sharding Introduction • Sharding (on Wikipedia) <p>To scale in, MongoDB has support for removing shards from your database.</p> <ul style="list-style-type: none"> • Removing Shards
<p>Multi-Datacenter Replication and Awareness</p>	<p>Oracle NoSQL Database supports DataCenters thru a non-electable replication group strategy. Read requests use nodes locally due to latency awareness in the client driver. Write availability is achieved in a local quorum though replicating to non-electable nodes in other data centers. This allows failures in a given data center to have no impact on read availability of the cluster as a whole, just possibly some reduced latency. Writes will always be performed at the currently elected master.</p>	<p>MongoDB can be configured to run in multiple datacenters via various options.</p> <ul style="list-style-type: none"> • Datacenter Awareness
<p>Graphical Monitoring/Admin Console</p>	<p>Oracle NoSQL Database provides proprietary, SNMP and JMX based protocols for monitorability of the cluster. The proprietary protocols are support thru both browser based and CLI interfaces. SNMP and JMX facilitate</p>	<p>MongoDB does not ship with a graphical monitoring/admin console. However, several community projects have developed graphical monitoring/admin programs.</p> <ul style="list-style-type: none"> • Monitoring and Diagnostics • Admin UIs

	<p>integration into monitoring systems like BMC and Ganglia.</p> <ul style="list-style-type: none">• Visual Admin Console• Standardized Monitoring Protocols• Command Line Admin	<p>10Gen offers a hosted monitoring service.</p>
--	--	--