Oracle NoSQL Database

Compared to Couchbase

Overview

- Oracle NoSQL Database is licensed under AGPL while Couchbase is Apache 2.0 licensed for its source code distribution, but it has proprietary license terms for binary distributions that are unique to Couchbase.
- 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.
- Couchbase is a memory based implementation of a document database with replication and journaling.

Comparison

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

Feature/Capability	Oracle NoSQL Database	Couchbase
Data Model	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 • Record Design Considerations • Avro Schemas	Couchbase is a JSON-based document datastore. Like other document datastores, records have no intrinsic relationships, and are stored in buckets. Value size is limited to 20Mbyte. • Store an Object?
Storage Model	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 representations, but must be moved at the granularity of these partitions. • BDB Storage - NoSQL before NoSQL was	Couchbase 2.0 is largely memory-based, asynchronously persisting data using a CouchDB fork and C library "couchstore" (prior versions of Couchbase use the SQLite storage engine). • Persistence • Couchbase Format

	cool • The evolution of BerkeleyDB	
Data Access and APIs	Oracle NoSQL Database has client library API's for Java and C. In the works are a Command Line Interface and Javascript API. • Client APIs	Couchbase provides drivers in several languages to access data through its binary memcached protocol. Couchbase also provides a REST API to monitor and manage a cluster (though it is not used to directly manage stored data). • Client Interface • Client-Libraries • Management REST API
Query Types and Query-ability	Oracle NoSQL Database provides key access methods (put, get, delete) including multi-key variations with large result set streaming support. The database can also be accessed using SQL as an external table from within a relational database. It is integrated with and can participate in MapReduce operations from a Hadoop environment. • Searching in Oracle NoSQL • External Table Support • NoSQL and MapReduce • Using Range Queries	Couchbase also provides four query options • ID lookups • MapReduce Views • UnQL Hadoop support is also possible through a plugin that streams data to a Hadoop Distributed File System (HDFS) or Hive for processing. • Hadoop Connector
Data Versioning and	Oracle NoSQL Database provides control at the	Couchbase is strongly consistent within a datacenter, replicating data between

Consistency	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. Flexible Consistency options	nodes in a cluster for failover. Inter-datacenter replication follows an eventually consistent CouchDB replication model. Via CouchDB, documents are internally revisioned (stored in a "_rev" value). However, prior revisions will be removed on a file compaction operation, making them unreliable. • Couchbase Architecture • Internal Version Field
Concurrency	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, all. Durability Guarantees	Couchbase claims to be ACID-compliant on a per-item basis, but has no multi-operation transactions. Couchbase clients connect to a server list (or via a proxy) where keys are sharded across the nodes. Couchbase nodes inherit memcached's default (and recommended) connection limit of 10k. • Transaction and concurrency • Cluster Design • Client-side Proxy
Replication	Oracle NoSQL Database supports replication for both availability and scalability. It	Couchbase supports two types of replication. For intra-datacenter clusters, Couchbase uses membase-

	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. • Replication configuration 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. • Topologies	style replication, which favors immediate consistency in the face of a network partition. For multi-datacenter deployments, CouchDB's master-master replication is used. • CouchDB Replication • Memcache Tap • CouchDB, Couchbase, Membase
Scaling Out and In	Oracle NoSQL Database scales out by redistribution of data partitions to newly added 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	Couchbase scales elastically by auto-sharding. They can be rebalanced to grow or shrink through the administrative interface. • Rebalancing • Clone to Grow with Auto Sharding

	administrator has the option of just letting it go or throttling or running during certain windows of time, pausing the process, etc. • Managing Topology Changes	
Multi-Datacenter Replication and Awareness	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.	Couchbase 2.0 supports cross-datacenter replication (XDCR). • Stabilizing Couchbase Server 2.0
Graphical Monitoring/Admin Console	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 integration into monitoring systems like BMC and Ganglia. • Visual Admin Console • Standardized	Couchbase provides a web-based monitoring/admin console. • Admin Wed Console

	Monitoring Protocols • Command Line Admin	