



www.couchbase.com

3250 Olcott Street, Santa Clara, CA 95954

Tel: +1 650-417-7500

Email: sales@couchbase.com

Couchbase Server

The company


Couchbase was formed in 2011 from the merger of Membase and CouchOne. It is VC backed and has its headquarters in California with additional offices in France, India and the United Kingdom. The company has a substantial partner programme.

Couchbase Server is a packaged version of its open source technology and is available in both community and commercial versions. Extending the NoSQL database to the edge, Couchbase Mobile runs on both mobile devices and other small footprint devices. For Couchbase

Server both on-premises and cloud-based deployments are supported. The company targets its database across industry boundaries and, prior to the most recent release, the focus has been on operational use cases. It has some notable clients, including Cisco, eBay, Sky, Telefonica, Ryanair, Tesco and many others.

What is it?

Couchbase Server is an open source, shared nothing, distributed, NoSQL document-oriented database that stores JSON documents. The product was originally developed with an architecture based on a distributed cache. A key-value store was then added and this has subsequently evolved into a multi-model, document-oriented design. It is important to understand this because Couchbase can still be deployed in any of these three ways. In 2015 the company introduced N1QL (pronounced Nickel), which is an extension to SQL. As one might expect, the company implemented a (rules-based) database optimiser at the same time. This was followed in 2017 with support for full text search. The current version, Couchbase Server 6.0, was released in November 2018 to support hybrid operational and analytic capabilities through the Couchbase Analytics service. Multi-dimensional scaling allows you to not only scale out by expanding the size of your cluster but also to scale each server node up on an SMP basis.



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

At the same time you can also allocate particular types of processing tasks (query processing, indexing and so on) to particular nodes within the cluster.

Couchbase is an append only database, which means that it does not delete old data when an update is performed. This has significant advantages in terms of performance. A facility known as compaction is used to ensure that storage requirements do not get out of hand. Apart from N1QL you can access Couchbase through a wide variety of languages, including Python and Scala and R via Couchbase ODBC/JDBC.

Couchbase primarily supports eventual consistency though immediate consistency is selectable on a per-operation basis. ACID support is also limited because atomicity is only supported within a document and not across documents.

How does it work?

Figure 1 illustrates the analytic architecture introduced in Couchbase Server 6.0. Some of this may need explanation. Specifically, internal Database Change Protocol (DCP), is a parallel capability that streams changes (updates and so forth) to all nodes, and keeps services synchronised, reducing and providing faster view

“The current version, Couchbase Server 6.0, was released in November 2018 to support hybrid operational and analytic capabilities through the Couchbase Analytics service.”

Architecture	★★★★★
Concurrent analytics	★★★★☆
Criticality	★★★★★
Ease of use	★★★★☆

Flexibility	★★★★★
In-process analytics	★★★★★
Performance	★★★★★
Scale	★★★★★

“ Couchbase offers all the advantages you would expect of a NoSQL database with schema-on-read, horizontal scalability and workload isolation along with the sort of high availability and resilience you would expect. ”

consistency. More generally, there is an embedded parallel engine to support performance for data intensive jobs. In addition, Couchbase Analytics supports parallel search, parallel joins (hash, index nested loops, and broadcast joins), parallel group-by (both pre-sorted and hash-based) and parallel sorts. All of these, incidentally, support spill-to-disk capabilities in case memory is insufficient.

There are several things that are not shown in this diagram. Notably, support for Bloom Filters, which are used to improve search efficiency; ODBC/JDBC connectivity and Spark integration. Jupyter Notebooks are supported through a Python SDK but the product does not support R or PMML (predictive modelling mark-up language). While the company is visualisation agnostic the lack of a pure SQL interface, integration with third party business intelligence tools is achieved through ODBC/JDBC drivers. The company has a partnership with Knowi. Finally, there is an important facility providing workload isolation that allows you to assign specific nodes in your cluster to complex analytics, indexing, ad hoc queries and so forth, as required, supporting the separation of operational and analytic workloads.

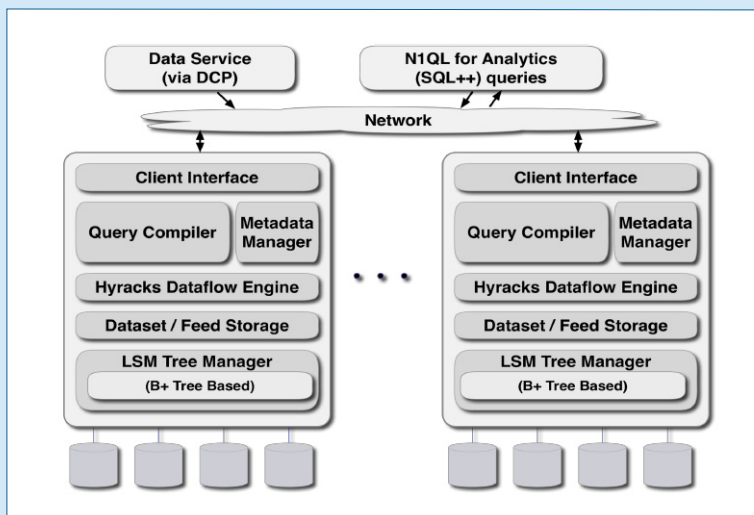


Figure 1 – The analytic architecture of Couchbase Server 6.0

Allied with this is N1QL for Analytics (and an Analytics Query Editor), which is based on SQL++. This is a language developed to bring SQL-like, declarative capabilities to environments where the data is semi-structured rather than structured, and N1QL for Analytics is the first commercial implementation of the SQL++ Framework. As can be seen, Couchbase Server uses log-structured-merge trees, which are good for high-speed ingestion, especially where indexed access (and Couchbase supports secondary indexes) is required.

Why should you care?

If we distinguish between “*transactions*” and “*operations*” by saying that the former require ACID guarantees, while the latter do not, then Couchbase Server with N1QL for Analytics is primarily targeted at hybrid environments that require the real-time processing of both operational data and analytics. As primarily a document database, Couchbase offers all the advantages you would expect of a NoSQL database with schema-on-read, horizontal scalability and workload isolation along with the sort of high availability and resilience you would expect of an enterprise-grade product.

The Bottom Line

The vast majority of environments are operational – as defined above – rather than transactional. If that is your use case and you want to combine operational processing with low latency analytics then Couchbase is well suited to these sorts of deployments.

FOR FURTHER INFORMATION AND RESEARCH [CLICK HERE](#)