



Couchbase

Why NoSQL?

Why successful enterprises
rely on NoSQL



Customer demand is driving Enterprises to adopt platforms powered by NoSQL, such as Couchbase's database

NoSQL is a common database type that stores information in **JSON documents** instead of the columns and rows used by **relational databases**. As a result, **NoSQL databases** are built from the ground up to be flexible, scalable, and capable of rapidly responding to the data management demands of modern, internet-based businesses.

Why enterprises choose NoSQL databases over relational databases

Customer experience has quickly become the most important competitive differentiator and businesses are working to meet these new expectations with services that are on-demand, real-time, resilient and responsive. Today's enterprises are interacting digitally - not only with their customers, but also with their employees, partners, vendors, and even their products - at an unprecedented scale. This interaction is powered by the internet and other 21st century technologies. At the heart of the revolution are a company's cloud, mobile, social media, Big Data, and IoT applications.

Most of this data is unstructured, amplifying the need of NoSQL's flexible, schema-less data model

Often, technical leaders are tasked with building an enterprise data management infrastructure that includes the following characteristics:

- Support large numbers of concurrent users (tens of thousands, perhaps millions)
- Deliver highly responsive experiences to a globally distributed base of users
- Be always available - no downtime
- Handle semi- and unstructured data
- Rapidly adapt to changing requirements with frequent updates and new features

SQL-based relational databases are unable to meet these new requirements whereas NoSQL databases can.

Consider just a few examples of Global 2000 enterprises that are deploying NoSQL for mission-critical applications that have been featured in recent news reports:

- **Tesco**, Europe's no. 1 retailer deploys NoSQL for e-commerce, product catalog, and other applications
- **Ryanair**, the world's busiest airline uses NoSQL to power its mobile app serving over 3 million users
- **Marriott** deploys NoSQL for its reservation system that books \$38 billion annually
- **Gannett**, the no. 1 U.S. newspaper publisher uses NoSQL for its proprietary content management system, Presto
- **GE** deploys NoSQL for its Predix platform to help manage the industrial internet



Five advantages of NoSQL databases

Ambitious customers with big ideas for how to use data have unleashed a new set of technology requirements for CIOs and technical leaders. Here are five trends that play into the advantages of NoSQL databases for addressing the challenges of operating in a digital economy.

Digital Economy Trends	Requirements
1. More customers are going online	<ul style="list-style-type: none">• Scaling to support thousands or even millions of users• Meeting UX requirements with consistent high performance• Maintaining availability 24 hours a day, 7 days a week
2. The internet is connecting everything	<ul style="list-style-type: none">• Supporting many different things with different data structures• Ensuring hardware and software are always updated with the latest security features• Supporting continuous streams of real-time data
3. Big data is getting bigger	<ul style="list-style-type: none">• Storing customer generated semi-structured and unstructured data• Storing different types of data from different sources in the same infrastructure or even the same cluster• Storing data generated by thousands or millions of customers and things
4. Applications are moving to the cloud	<ul style="list-style-type: none">• Scaling on demand to support more customers, store more data• Operating applications on a global scale to support customers worldwide• Minimizing infrastructure costs, achieving a faster time to market
5. The world has gone mobile	<ul style="list-style-type: none">• Creating “offline first” apps - network connection not required• Synchronizing mobile data with remote databases in the cloud• Supporting multiple mobile platforms with a single backend



Why relational databases fall short

Relational databases were born in the era of mainframes and business applications – long before the Internet, the cloud, Big Data, mobile and now, the Digital Economy. In fact, the first commercial implementation was released by Oracle in 1979. These databases were engineered to run on a single server – the bigger, the better. The only way to increase the capacity of these databases was to upgrade the servers – processors, memory, and storage – to scale up.

NoSQL databases emerged as a result of the exponential growth of the internet and the rise of web applications. Google released the BigTable research in 2006, and Amazon released the Dynamo research paper in 2007. These databases were engineered to meet a new generation of enterprise requirements:

The need to **develop with agility** and to **operate at any scale**.

Develop with agility

To remain competitive in the Digital Economy, enterprises must innovate – and now they have to do it faster than ever before. As this innovation centers on the development of modern web, mobile, and IoT applications, developers are under extraordinary pressure. Speed is critical, but so is agility, since these applications evolve far more rapidly than legacy applications like ERP. Relational databases require a relatively restricted flat data structure and don't respond well to frequent changes in the data model. This simply doesn't meet the needs of modern, frequently changing applications and business requirements.

Scoping for changing requirements

A core principle of agile development is adapting to evolving application requirements: when the requirements change, the data model also changes. This is a problem for relational databases because the data model is fixed and defined by a static schema. So in order to change the data model, developers have to modify the schema, or worse, request a “schema change” from the database administrators. This slows down or stops development, not only because it is a manual, time-consuming process, but it also impacts other applications and services.



Flexibility for faster development

By comparison, a NoSQL document database fully supports agile development, because it is schema-less and does not statically define how the data must be modeled. Instead, it defers to the applications and services, and thus to the developers as to how data should be modeled. With NoSQL, the data model is defined by the application model. Applications and services model data as objects.



Simplicity for easier development

Applications and services model data as objects (e.g. employee), multi-valued data as collections (e.g., roles), and related data as nested objects or collections (e.g. manager). However, relational databases model data as tables of rows and columns - related data as rows within different tables, multi-valued data as rows within the same table. The problem with relational databases is that data is read and written by disassembling, or “shredding,” and reassembling objects. This is the object-relational “impedance mismatch.” The workaround is object-relational mapping frameworks, which are inefficient at best, problematic at worst.

As an example, consider an application for managing resumes. It interacts with resumes as an object, the user object. It contains an array for skills and a collection for positions. However, writing a resume to a relational database requires the application to “shred” the user object.

Storing this resume would require the application to insert six rows into three tables, as illustrated in **Figure 3**.

However, reading this profile would require the application to read six rows from three tables, as illustrated in **Figure 4**.



Figure 4

RDBMS - Queries return duplicate data, applications have to filter it out.

Shane	Johnson	Big Data	Product Marketing	Couchbase
Shane	Johnson	Big Data	Technical Marketing	Red Hat
Shane	Johnson	Java	Product Marketing	Couchbase
Shane	Johnson	Java	Technical Marketing	Red Hat
Shane	Johnson	NoSQL	Product Marketing	Couchbase
Shane	Johnson	NoSQL	Technical Marketing	Red Hat

In contrast, a document-oriented NoSQL database reads and writes data formatted in JSON - which is the de facto standard for consuming and producing data for web, mobile, and IoT applications. It not only eliminates the object-relational impedance mismatch, it eliminates the overhead of ORM frameworks and simplifies application development because objects are read and written without "shredding" them - i.e., a single object can be read or written as a single document, as illustrated in **Figure 5**.



What about querying and SQL?

The Couchbase NoSQL Server supports N1QL, an extension of SQL that supports JSON documents. It not only supports standard SELECT / FROM / WHERE statements, it also supports aggregation (GROUP BY), sorting (SORT BY), joins (LEFT OUTER / INNER), as well as querying nested arrays and collections. In addition, query performance can be improved with composite, partial, covering indexes, and more.



SQL	NIQL
<pre>SELECT breweries.name AS brewery, count(*) AS cnt FROM beers INNER JOIN breweries ON beer.brewery_id = breweries.id WHERE beers.type = "beer" AND breweries.type = "brewery" AND beers.style = "American-Style Imperial Stout" GROUP BY breweries.name HAVING count(*) > 2 ORDER BY cnt DESC;</pre>	<pre>SELECT breweries.name AS brewery, count(*) AS cnt FROM `beer-sample` beers INNER JOIN `beer-sample` breweries ON KEYS beers.brewery_id WHERE beers.type = "beer" AND breweries.type = "brewery" AND beers.style = "American-Style Imperial Stout" GROUP BY breweries.name HAVING count(*) > 2 ORDER BY cnt DESC;</pre>

What about ACID Transactions?

When you flatten out a business entity into multiple separate tables, you require a transaction for almost every update. With NoSQL databases, you don't need to flatten out the entity, but can usually contain it in a single document. Updates to a single document are atomic and don't require a transaction. However, there are updates that span entities but require all-or-nothing semantics. This is why NoSQL databases like Couchbase now support transactions.

The addition of transactions and SQL greatly expand the number of use cases where a NoSQL database can be considered. In the past, the inability to join or handle transactional operations meant that NoSQL databases were only chosen for the highest volume and scale use cases. The addition of SQL and Transactions means that NoSQL databases can also be chosen for traditional RDBMS cases where volume is lower but the need for flexibility is higher. Additionally, by selecting a transactional NoSQL database, many traditionally complex applications can be simplified because there is no need for an object-relational-mapping (ORM) tool.

Operate at any scale

Databases that support web, mobile, and IoT applications must be able to operate at any scale. While it is possible to scale a relational database like Oracle (using, for example, Oracle RAC), doing so is typically complex, expensive, and not fully reliable. With Oracle, for example, scaling out using RAC technology requires numerous components and creates a single point of failure that jeopardizes availability. By comparison, a NoSQL distributed database – designed with a scale-out architecture and no single point of failure – provides compelling operational advantages.


Elasticity for Performance at scale

Applications and services have to support an ever-increasing number of users and data – hundreds to thousands to millions of users, and gigabytes to terabytes of operational data. At the same time, they have to scale to maintain performance, and they have to do it efficiently.




The database has to be able to scale reads, writes, and storage.

This is a problem for relational databases that are limited to scaling up – i.e., adding more processors, memory, and storage to a single physical server. As a result, the ability to scale efficiently, and on demand, is a challenge. It becomes increasingly expensive, because enterprises have to purchase bigger and bigger servers to accommodate more users and more data. In addition, it can result in downtime if the database has to be taken offline to perform hardware upgrades.



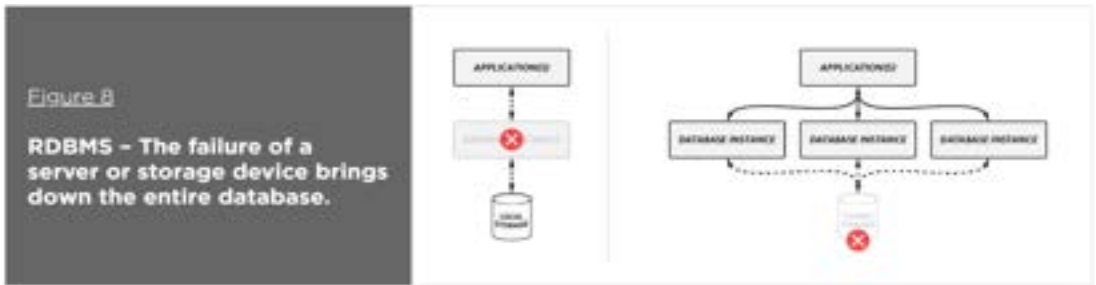
A distributed NoSQL database, however, leverages commodity hardware to scale out – i.e., add more resources simply by adding more servers. The ability to scale out enables enterprises to scale more efficiently by (a) deploying no more hardware than is required to meet the current load; (b) leveraging less expensive hardware and/or cloud infrastructure; and (c) scaling on-demand and without downtime.



By distributing reads, writes and storage, NoSQL databases are able to operate at any scale. Additionally, they were designed to be easy to configure, install and manage both small and large clusters.

Availability for always-on, global deployment

As more and more customer engagements take place online via web and mobile apps, availability becomes a major, if not primary, concern. These mission-critical applications have to be available 24 hours a day, 7 days a week – no exceptions. Delivering 24x7 availability is a challenge for relational databases that are deployed to a single physical server or that rely on clustering with shared storage. If deployed as a single server and it fails, or as a cluster and the shared storage fails, the database becomes unavailable.



In contrast to relational technology, a distributed, NoSQL database partitions and distributes data to multiple database instances with no shared resources. In addition, the data can be replicated to one or more instances for high availability (intercluster replication). While relational databases like Oracle require separate software for replication, for example, Oracle Active Data Guard, NoSQL databases do not - it's built in and it's automatic. In addition, automatic failover ensures that if a node fails, the database can continue to perform reads and writes by sending the requests to a different node.



Customer behavior now requires organizations to support multiple physical, online and mobile channels in multiple regions and often multiple countries. While deploying a database to multiple data centers increases availability and helps with disaster recovery, it also has the benefit of increasing performance too. All reads and writes can be executed on the nearest data center, thereby reducing latency.

Ensuring global availability is difficult for relational databases where separate add-ons are required - which increase complexity - or where replication between multiple data centers can only be used for failover, because only one data center is active at a time. Oracle, for example, requires Oracle GoldenGate. When replicating between data centers, applications built on relational databases can experience performance degradation or find that the data centers are severely out of sync.



A distributed, NoSQL database includes built-in replication between data centers – no separate software is required. In addition, some include bidirectional replication enabling full active-active deployments to multiple data centers. This enables the database to be deployed in multiple countries or regions while providing local data access to local applications and their users. Deploying to multiple data centers not only improves performance, but enables immediate failover via hardware routers. Applications don't have to wait for the database to discover the failure and perform its own failover.



NoSQL is a better fit for Digital Economy requirements

As enterprises shift to the Digital Economy – enabled by cloud, mobile, social media, and big data technologies – developers and operations teams have to build and maintain web, mobile, and Internet of Things (IoT) applications faster and faster, and at a greater scale. NoSQL is increasingly the preferred database technology to power today's web, mobile, and IoT applications.

Hundreds of Global 2000 enterprises, along with tens of thousands of smaller businesses and startups, have adopted NoSQL. For many, the use of NoSQL started with an open source cache, proof of concept or a small application, then expanded to targeted mission-critical applications, and is now the foundation for all application development. Today, the Couchbase DB serves thousands of these types of customers.

With NoSQL, enterprises are better able to both develop with agility and operate at any scale – and to deliver the performance and availability required to meet the demands of Digital Economy businesses.

About Couchbase

Unlike other NoSQL databases, Couchbase provides an enterprise-class, multicloud to edge database that offers the robust capabilities required for business-critical applications on a highly scalable and available platform. As a distributed cloud-native database, Couchbase runs in modern dynamic environments and on any cloud, either customer-managed or fully managed as-a-service. Couchbase is built on open standards, combining the best of NoSQL with the power and familiarity of SQL, to simplify the transition from mainframe and relational databases.

Couchbase has become pervasive in our everyday lives; our customers include industry leaders Amadeus, American Express, Carrefour, Cisco, Comcast/Sky, Disney, eBay, LinkedIn, Marriott, Tesco, Tommy Hilfiger, United, Verizon, as well as hundreds of other household names. For more information, visit www.couchbase.com.

© 2020 Couchbase. All rights reserved.

