



Relational to NoSQL: Getting Started From SQL Server



WHY THE SHIFT TO NOSQL?

*

NOSQL HAS BECOME A FOUNDATION FOR MODERN WEB, MOBILE, AND IOT APPLICATION DEVELOPMENT. AT COUCHBASE, WE'VE ENABLED HUNDREDS OF ENTERPRISES, AS WELL AS NUMEROUS GROWTH COMPANIES AND STARTUPS, TO DEPLOY NOSQL FOR BETTER AGILITY, PERFORMANCE, AND LOWER COSTS As enterprises modernize, teams have to build and maintain applications more rapidly and at greater scale. Applications must be resilient and available whether their clients are web, mobile, or the Internet of Things (IoT). If any of these channels fail, customers go elsewhere. Today, enterprises in every industry from travel, to technology, to retail and services are leveraging NoSQL database technology for more agile development, reduced operational costs, and scalable operations. For many, the use of NoSQL started with a cache, proof of concept (POC), or small application, then expanded to targeted mission-critical applications. NoSQL has now become a foundation for modern web, mobile, and IoT application development.

Some of the largest internet and enterprise companies are using NoSQL technology to deploy their mission-critical applications. Examples include:

- **Gannett** publisher of USA Today and 90+ media properties, replaced relational database technology with NoSQL to power its digital publishing platform
- **Marriott** deployed NoSQL to modernize its hotel reservation system that supports \$38 billion in annual bookings
- **FHL Bank Topeka** integrates NoSQL with SQL Server to speed up access to customer financial data for its 770 member banks
- **Cars.com** with over 30 million visits per month, replaced SQL Server with NoSQL to store customer and vehicle data

Couchbase has enabled hundreds of enterprises, growth companies, and startups to deploy NoSQL for better flexibility, fast performance, and affordable costs. The goal of this paper is to help you introduce NoSQL into your organization by highlighting lessons learned from teams that successfully adopted NoSQL. We'll explore key considerations and strategies for transitioning to NoSQL, in particular, to a document database (Couchbase Server or Couchbase Capella[™] DBaaS), with tips for moving from SQL Server and other relational databases. Note, there are use cases in which NoSQL is not a replacement for, but a complement to, existing infrastructure, facilitating the use of polyglot persistence.

We'll start with recommendations for identifying and selecting the right application. Next, we'll cover strategies for modeling relational data as documents, how to access them within your application, and how to migrate data from a relational database. Finally, we'll highlight the basics of operating a NoSQL database in comparison to a relational database.

TOP 5 REASONS COMPANIES REPLACE SQL SERVER WITH A NOSQL DATABASE

If you're running into limits with Microsoft SQL Server (or other relational databases) – either in terms of rising costs and complexity, or in scaling to meet your requirements – this is the time to evaluate a NoSQL database.



Many companies have chosen Couchbase Server either to augment Microsoft SQL Server and other relational databases, or in some cases replace them.

Their top 5 reasons:

1. Flexibility to scale

Whether running on three nodes or hundreds – in the cloud or on-prem – Couchbase Server is based on a distributed architecture to scale on commodity hardware with ease, regardless of the operating system. Couchbase Capella provides a fully hosted, managed cloud DBaaS to make scaling and operations even easier.

2. Faster performance

Couchbase Server features an integrated cache, memory-optimized indexes, and memory to-memory replication to deliver consistent, high throughput with submillisecond response time, at any scale.

3. Up to 40x more affordable

With subscriptions based on the number of instances, Couchbase typically costs 5-40x less than SQL Server and its per-core based licensing – especially on servers with many cores.

4. Familiarity

Couchbase fuses the best of relational with the best of NoSQL. Data is stored as flexible JSON, but can be queried with SQL++. With ACID transactions also possible, your team will have a running start in modernizing.

5. Future-proof

Couchbase supports the growth of use cases. While you may start using Couchbase for its high-performance caching capabilities, as your application expands, you can use built-in query, text search, mobile, analytics, syncing, and more as needed.

WHY THE SHIFT TO NOSQL?

Many enterprises have successfully introduced NoSQL by identifying a single application or service to start with. It could be a new one that is being developed, or an existing application that's being refactored. Examples include:

- A high performance, highly available caching service
- A small, independent application (or microservice) with a narrow scope
- A logical or physical service within a large application
- A global service that powers multiple applications across multiple regions







MANY ENTERPRISES HAVE SUCCESSFULLY INTRODUCED NOSQL BY IDENTIFYING A SINGLE APPLICATION OR SERVICE TO START WITH.

IDEAL CANDIDATES HAVE ONE OR MORE OF THE FOLLOWING CHARACTERISTICS OR REQUIREMENTS:

- Need to modify the data model to respond to change frequently e.g., to accommodate new preferences, new social media accounts, etc.)
- Ability to read and write JSON documents (semi-structured data) to/from web and mobile clients
- Provide low latency, high throughput access to data e.g., users expect interactions to be instant, and waiting negatively impacts the experience
- Support thousands to millions of concurrent users e.g., the number of users rapidly increases, sometimes exponentially, as when content goes viral
- Support users in any country or region they're everywhere
- Support users and be available 24x7
- Store terabytes of data
- Deploy in multiple data centers with an active/active configuration
- Build enterprise applications a series of applications with constantly changing schema requirements, complex or unstructured data

SOME COMMON EXAMPLES OF GOOD USE CASES FOR NOSQL:

- Product catalog service
- Asset tracking service
- Content management service
- Application configuration service
- Customer management service
- File or streaming metadata service

MODELING AND MIGRATING YOUR DATA

Couchbase Server is a **document database** – data is stored in JSON document collections, instead of tables. While relational databases rely on an explicit predefined schema to describe the structure of data, document databases do not – JSON documents are self-describing. As such, every JSON document includes its own flexible schema, and it can be changed on demand by changing the document itself.



RELATIONAL DATABASE



Column	Format
first_name	VARCHAR
last_name	VARCHAR
birth_date	DATE

	OCUMENT DATABASE		
{	"firstName": "Shane", "lastName": "Johnson", "birthDate":"08/12/1979"		

TABLE / ROW:

first_name	last_name	birth_date
Shane	Johnson	08/12/1979

Figure 1: Relational schema and data vs. self-describing JSON documents

It's important to understand that JSON documents are not limited to primitive fields. They can include arrays and objects, and they can be nested, just like applications. For this reason, there is no "impedance mismatch" between application objects and JSON documents. No complex object-relational mapping (ORM) solution is required.



Figure 2: Multiple tables vs. nested data with JSON documents

Just as every row in a table requires a primary key, every document requires an object ID. Many applications rely on relational databases to automatically generate the primary key (for example, with the IDENTITY columns in Microsoft SQL Server). Document databases can use unique keys like UUID/GUID, but applications can also use natural keys where possible.



COUCHBASE SERVER IS A DOCUMENT DATABASE – DATA IS STORED IN JSON DOCUMENTS, NOT IN TABLES.

EVERY JSON DOCUMENT INCLUDES ITS OWN FLEXIBLE SCHEMA, AND IT CAN BE CHANGED ON DEMAND BY CHANGING THE DOCUMENT ITSELF.



In a relational database, primary keys are defined per table. It's not uncommon for the primary key of different rows in different tables to have the same value. After all, a row is identified by its table and primary key. However, document databases do not store documents in tables; in Couchbase Server, they're stored in collections (which are then stored in scopes and then buckets). You can store any type of document within a collection, but typically you will store similar types of documents within the same collection (i.e., very similar to using a relational table).



Figure 3: Couchbase Server – Documents are stored in buckets, buckets are partitioned and ditributed aross nodes automatically.

The benefit of using natural keys with a document database is that a document can be identified by an object ID, even if the collection stores different types of documents.

For example, consider a single collection with blogs, authors, and comments stored in separate documents:

author::shane

author::shane::blogs

blog::nosql_fueled_hadoop

blog::nosql_fueled_hadoop::comments

These object IDs not only enable the bucket to store related documents; they're also human readable, deterministic, and semantic. In addition, an application can construct these keys easily to fetch or query using them. Even if you decided to put these four documents into four separate buckets, these keys will still be beneficial.

A document can be modeled after a row (flat), or it can be modeled after related rows in multiple tables (nested). However, documents should be modeled based on how applications interact with the data. While some documents may contain nested data, others may reference it.





In the absence of tables, applications can benefit from including a field in documents that identifies the type of document. In the figure 4 example, there are user, address, and account types. The type field can then be indexed to improve query performance when multiple types of documents are stored within the same bucket.

In relational databases, children reference their parents via foreign keys. However, in document databases, parents are able to reference their children when appropriate. That's because, while a field in a row can only contain a single value, a field within a document can contain multiple values. In the example of a related model (figure4), the addresses and accounts fields contain multiple object IDs – one for each address. However, the shipping and billing object IDs are not required – they are deterministic.

A document can contain fields, objects, lists, and arrays. In the examples above, the addresses field contains a list of addresses or references to address. The accounts field contains an array of accounts or references to accounts.



NOTE: **Couchbase Server** supports atomic counters, and like IDENITY columns in Microsoft SQL Server, they can be used to automatically generate object IDs for documents. In addition, counters can be incremented by any amount. Like IDENTITY columns in Microsoft SQL Server, counters can be incremented by any number, for example, 50.

MIGRATION TIP: Create object IDs for documents that include the row's primary key. For example, if the primary key of a row in the products table is 123, the document ID is product::123. However, if you are putting documents in a product collection, you may want to omit "product::" to save memory and disk space.

REFERENCE OR NEST RELATED DATA?

There are two things to consider when deciding how to model related data:

1. Is it a one-to-one or one-to-many relationship?

2. How often is the data accessed?

If it's a one-to-one or one-to-many relationship (a child has one parent), it may be better to store the related data as nested objects. This approach results in a simple data model and reduces or eliminates the need to query multiple documents However, if it's a many-to-one or many-to-many relationship (a child has multiple parents), it may be better to store the related data as separate documents, which reduces or eliminates the need to maintain duplicate data.

If a majority of the reads are limited to parent data (e.g., first and last name), it may be better to model the children (e.g., addresses and accounts) as separate documents. This results in better performance, because the data can be read with a single key-value operation instead of a query, and reduces bandwidth, because the amount of data being transferred is smaller. However, if a majority of the reads include both parent and child data, it may be better to model the children as nested objects. This approach results in great performance because the data can be read with a single key-value operation instead of a query.

If a majority of the writes are to the parent or child, but not both, it may be better to model the children as separate documents. For example, if user profiles are created with a wizard – first add info, then add addresses, finally add accounts – or if a user can update an address or account without updating their info. However, if a majority of writes are to parent and child (both) – for example, there's a single form to create or update a user – it may be better to model the children as nested objects.

When to nest? Considerations:

- One-to-one or one-to-many? Nest.
- Most reads are for parent and child together? Nest.
- Many-to-one or many-to-many? Don't nest.
- Most writes are for parent or child? Don't nest.
- Most reads are for parent data? Don't nest.
- Most writes are for parent and child together? Nest.





Finally, it may be better to model children as separate documents to reduce document size and write contention. For example, the number of reviews on a product may grow indefinitely. If they were embedded, the size of the product document could become excessive, resulting in slower reads. Consider a blog and comments. When a blog is first published, there may be a lot of readers posting comments. If the comments are embedded, many concurrent users will try to update the same blog document at the same time, resulting in slower writes. A good compromise may be to store comments as separate threads – a document for every top-level comment that embeds all replies.

<pre>EPARATE BILLING AND SHIPPING ADDRESSES { "docType": "user", "firstName": "Shane", "lastName": "Johnson, "addresses": { "billing": { "city": "Chicago", "state": "IL", }. </pre>	ONE ADDRESS FOR BILLING AND SHIPPING { "docType": "user", "firstName": "Brendan", "lastName": "Bond, "addresses": { "billing": { "city": "Half Moon Bay", "state": "CA", }.
"shipping": { "city": "San Jose", "state": "CA", } }	"shipping": "same"} }

Figure 5: Different documents of the same type can have different schema



PERFORMING A MIGRATION?



NOSQL DATABASES PROVIDE DATA ACCESS VIA KEY-VALUE APIS, SQL++, FULL-TEXT SEARCH, AND MORE. **NOTE:** SAME OBJECT, DIFFERENT STRUCTURE? AS ILLUSTRATED IN FIGURE 5, IT'S POSSIBLE FOR THE SAME FIELD OR OBJECT TO HAVE A DIFFERENT STRUCTURE IN DIFFERENT DOCUMENTS.

)

The easiest and fastest way to get started is to export your relational data to CSV files, and import them into Couchbase Server. This may not represent the final data model, but it will enable you to start interacting with Couchbase Server right away. Couchbase Server includes a command-line utility, cbimport, for importing data in CSV files.

\$ cbimport csv -c couchbase://127.0.0.1 -u Administrator -p
password -b default --scope collection-exp myscope.products -d
file:///products.csv -g %id% -t 4

UNDERSTANDING YOUR ACCESS PATTERNS

NoSQL databases provide data access via key-value APIs, SQL++ query APIs, or SDKs. The key-value API provides the best performance – since it will often be direct from an in-memory cache. The query API or language provides the most power and flexibility – enabling applications to sort, filter, transform, group, and combine documents. Queries to Couchbase are done using SQL, just like in a relational database (with extensions for JSON, hence "SQL++").

KEY VALUE

The key-value API can provide a great deal of data access without the need to perform queries. In the example below, once you have the object ID of the user profile document, you can figure out what the object IDs of the address and account documents are.

OPTION ONE: RELATED DATA IS REFERENCED



2
3 bucket.get("user::100::address::billing");
4
5 bucket.replace("user::100::address::shipping", address);
6
7 bucket.insert("user::100::account::002, account);

bucket.get("user::100");

12

8

OPTION TWO: RELATED DATA IS NESTED



+

COUCHBASE SERVER'S IMPLEMENTATION OF THE SQL++ STANDARD IS CALLED N1QL, WHICH EXTENDS SQL TO JSON DOCUMENTS.

QUERY

The query API or language, combined with proper indexing, can provide a great deal of power and flexibility without sacrificing performance. Couchbase Server provides a SQL++ implementation called N1QL, which extends SQL to JSON documents. N1QL also has support for ANSI joins making it easier for developers to apply their SQL knowledge to develop applications within Couchbase.



et all Platinum users	1	SELECT	firstName, lastName
	2	FROM	users
	3	WHERE	status = "Platinum";
	4		
Get all users with a Visa	5	SELECT	firstName, lastName
or Mastercard account	6	FROM	users u INNER JOIN accounts a
	7	ON KEYS	u.accounts
	8	WHERE	a.type = "Visa" OR
	9		a.type = "Mastercard";
	10		
Get all users with a billing	11	SELECT	firstName, lastName
address in California	12	FROM	users u INNER JOIN addresses a
	13	ON KEYS	u.addresses.shipping
	14	WHERE	a.state = "CA";
	15		



While one of the benefits of storing related data as separate documents is the ability to read a subset of the data (e.g., shipping address), the same thing can be accomplished with a query API or language when related data is nested. For example, to read the billing address from a user profile document that stores all related data as nested objects.

1	SELECT	addresses.billing	
2	FROM	users	
3	USE KEYS	["user::100"];	
4			



In addition, while one of the benefits of storing related data as nested objects is the ability to access all data with a single read, the same thing can be accomplished with a query API or language when related data is stored as separate documents. For example, to read the user profile and accounts and addresses when they are stored as separate documents.

1	SELECT	*
2	FROM	users
3	NEST	accounts ON KEYS users.accounts
4	NEST	shippingAddress
5	ON KEYS	users.addresses.shipping
6	NEST	billingAddress
7	ON KEYS	users.addresses.billing
8	USE KEYS	["user::100"];
9		

The query language can be used to perform CRUD operations as an alternative to the key-value API. This enables applications built on top of a relational database to migrate all data access by replacing SQL statements with SQL++ statements. One of the advantages of performing CRUD operations with SQL++ have the ability to perform partial updates:



N1QL abstracts the data model from the application model. Regardless of how data is modeled in the database, applications can query it any way they need to by joining documents, nesting and unnesting them, and more. It provides developers with the flexibility they need to model data one way and query it in many.

INDEXING YOUR DATA



NOSQL DATABASES SUPPORT INDEXES TO VARYING DEGREES - COUCHBASE SERVER INCLUDES COMPREHENSIVE INDEXING SUPPORT.



Query performance can be improved by indexing data. NoSQL databases support indexes to varying degrees – Couchbase Serverincludes comprehensive indexing support. Below are some indexing examples.

A **simple index** on the user status:

11	CREATE IND	EX user status	
2	ON	users(status).	
2		users(status),	
JERY			
JERY			
	SELECT on		
	SELECT co	unt(status)	
JERY	SELECT co FROM use	unt(status) ers	

A **composite index** on user status and shipping state:

1	CREATE INDEX user_status_state
2	ON users(status, addresses.shipping.state)
3	- Anna a'
JER	Y
UER	SELECT JactName firstName
	SELECT lastName, firstName
1 2	Y SELECT lastName, firstName FROM users
I 2 3	Y SELECT lastName, firstName FROM users WHERE status = "Platinum" AND

A functional index on shipping state:

1	CREATE INDEX	user_accounts
1	ON	users(count(accounts));
R	Y	
ER	Y SELECT lastNa	me, firstName
	Y SELECT lastNa FROM users	me, firstName



A **partial index** on user billing state of users with a Visa credit card:



Couchbase Server supports index intersection. A query can scan multiple indexes in parallel. As a result, it may not be necessary to create multiple indexes that include the same field, thereby reducing both disk and memory usage. You can also index a large number of documents and horizontally scale out an index as needed. The system will transparently partition the index **across a number of index nodes** using hash partitioning and will increase the performance and data capacity of the cluster.

CONNECTING TO THE DATABASE

Applications access data in NoSQL databases via clients. Couchbase Server SDKs are all topology-aware clients (e.g., smart clients) available in many languages: Java, Node.js, PHP, Python, C, and more. These clients are configured in much the same way JDBC/ODBC drivers are configured.

LATIONAL	COUCHBASE SERVER
<pre>1 SQLServerDataSource ds = new SQLServerDataSource();</pre>	1 List <string> nodes = Arrays.asList("127.0.0.1");</string>
 ds.setServerName("localhost"); ds.setPortNumber(1433); 	 CouchbaseCluster cluster = CouchbaseCluster.create(nodes);
5 6 Connection conn = ds.getConnection();	5 6 Bucket bucket = cluster.getBucket("users");
7 8 // access data via a connection	7 8 // access data via a bucket

Figure 6: Creating a connection to a relational database vs. Couchbase Server

A bucket is a higher-level abstraction than a connection, and a cluster can contain multiple buckets. In the example above, the application can access data in the users bucket. However, while key-value operations are limited to the users bucket, SQL++ queries are not.

Couchbase Server is a distributed database, but applications do not have to pass in the IP address of every node. However, they should provide more than one IP address so that if the first node is unavailable or unreachable, they can try to connect to the next node. After the client connects to a node, it will retrieve the IP address of the remaining nodes.

Couchbase Server clients also maintain a cluster map, which enables them to communicate directly with nodes. In addition, the cluster map enables operations teams to scale out the database without impacting the application. Regardless of the number of nodes, the application sees a single database. There are no application changes required to scale from a single node to dozens – the clients are automatically updated.

In addition, with Couchbase Server, applications no longer have to rely on objectrelational mapping frameworks for data access, because there is no impedance mismatch between the data model and the object model. In fact, domain objects are optional. Applications can interact with the data via document objects or by serializing domain objects to and from JSON.



APPLICATIONS ACCESS DATA IN NOSQL DATABASES VIA CLIENTS. COUCHBASE SERVER'S TOPOLOGY-AWARE SDKS ARE AVAILABLE IN MANY LANGUAGES: JAVA, NODE.JS, .NET, PHP, PYTHON, C, AND MORE. **NOTE:** IN ADDITION TO THE CLIENTS, THERE ARE SUPPORTED, CERTIFIED JDBC/ODBC DATABASE DRIVERS AVAILABLE FOR COUCHBASE SERVER.

Figure 7: working with domain objects vs document objects

It's easy to serialize domain objects to and from JSON, and it may be helpful to do so for applications with a complex domain model or business logic. However, for new applications or services, working with document objects will require less code and provide more flexibility – developers can change the data model without having to change the application model. For example, you can add a new field to a form without changing application code.

TRANSACTIONS

Similar to a relational database, Couchbase offers ACID transaction support. Transactions must be initiated by the Couchbase SDK, and can include key-value operations and/or SQL operations.

Transactions are not required when changing data in a single document structure as changes within a document are atomic. When architecting new data structures this should be taken into account as atomic changes to a single document are lighter weight.

```
try
{
    await transactions.RunAsync(async (ctx)=>
    {
        // 'ctx' is an AttemptContext, which permits getting, inserting,
        // removing and replacing documents, along with committing and
        // rolling back the transaction.
        // ... Your transaction logic here ...
        // This call is optional - if you leave it off, the transaction
        // will be committed anyway.
        await ctx.CommitAsync().ConfigureAwait(false);
}).ConfigureAwait(false);
}
catch (TransactionCommitAmbiguousException e)
{
    // The application will of course want to use its own logging rather
    // than Console.WriteLine
   Console.Error.WriteLine("Transaction possibly committed");
   Console.Error.WriteLine(e);
}
catch (TransactionFailedException e)
{
   Console.Error.WriteLine("Transaction did not reach commit point");
    Console.Error.WriteLine(e);
```

Figure 8: Transactions using the .NET API

One of the key advantages driving the adoption of NoSQL databases with a distributed architecture is their ability to scale faster, easier, and at a significantly lower cost than relational databases. While most relational databases are capable of clustering (e.g., Microsoft SQL Server), they are still limited to scaling up – failover clustering relies on shared storage while Always-on availability groups are limited to replication. As a result, more data requires a bigger disk, and more users require a bigger server. The shared storage not only becomes a bottleneck, it becomes a single point of failure. In contrast, most NoSQL databases are distributed to scale out – more data requires more disks, not a bigger one, and more users require more servers, not a bigger one.





Couchbase Server's topology-aware clients and consistent hashing distribute data within a cluster automatically. Data can be replicated to one or more nodes to provide high availability, also automatically.

INSTALLING COUCHBASE SERVER

Installing Couchbase Server requires little more than downloading the install binary, running it, and configuring the database via the web-based administrative console or CLI. It can also be installed via Docker.

Perhaps the easiest way to get started is with Couchbase's Database-as-a-Service (DBaaS), Couchbase Capella™.

COUCHBASE DELIVERED AS A SERVICE

With Capella, database management (setup, ongoing operations, and maintenance) is automated and streamlined so customers can focus on other areas like application development and improving time to market.



ONE OF THE KEY ADVANTAGES DRIVING ADOPTION OF NOSQL DATABASES WITH A DISTRIBUTED ARCHITECTURE IS THEIR ABILITY TO SCALE FASTER, EASIER, AND AT SIGNIFICANTLY LOWER COST THAN RELATIONAL DATABASES. NOSQL DATABASES SHOULD BE DEPLOYED AS A CLUSTER TO REALIZE ALL OF THEIR BENEFITS.

Key benefits include:

Fully managed

• Automated setup, backups, upgrades, and ongoing management to deliver an always-on service, reducingyour operational efforts.

Automated scaling in/out and up/down

• Easily add, remove, or change nodes to meet your current needs. Couchbase rebalances your data. No application changes needed.

Single pane for multi-cluster, multi-region

• The Capella control plane manages data across clusters and across clouds, allowing you to be cloud-provideragnostic. It also provides tools for SQL++ access, document viewing, index creation, and full-text search within the UI.

High availability

• Capella guarantees the global reliability of your data throughout regions and availability zones via native replication across geo-aware clusters, all day, every day.

Self-monitoring, self-healing

• Capella proactively monitors clusters 24/7 to locate, assess, and resolve issues automatically.

Security peace of mind

• In addition to SOC II compliance, Capella delivers end-to-end encryption from the SDK to the disk and offers granular role-based access control.



Figure 9: Topology-aware clients in Couchbase Server





MONITORING AND MANAGING YOUR DEPLOYMENT

Couchbase Server and Couchbase Capella include an integrated, comprehensive administration console as well as REST and CLI APIs.

While many relational and NoSQL databases require separate administration tools, Couchbase Server includes an integrated, comprehensive administration console as well as REST and CLI APIs.

The administration console and the REST/CLI APIs enable administrators to manage and monitor clusters, both small and large, with minimal effort. Functionality enabled through the Couchbase admin console includes:

Management, monitoring, and configuration

- Cluster/node/bucket/views
- Cross data center replication (XDCR)
- Database performance, network utilization, resource utilization

Tasks

- Add and remove nodes
- Failover nodes
- Rebalance cluster

Configuration

- Authentication and authorization
- Auditing

Authentication and authorization

- Audit
- Monitor
- View and collect log information

In addition to the administration console and APIs, Couchbase Server includes a number of command line tools to perform additional tasks such as, among others:

- · cbbackupmgr Full, cumulative, and incremental backup and restore
- cbcollect_info and cbdstats Gather node and cluster diagnostics (280+ metrics)
- cbq Run SQL++ queries from the command line
- · cbimport/cbexport Transfer data to and from JSON or CSV files

PUTTING IT ALL TOGETHER: HOW TO CONDUCT A SUCCESSFUL PROOF OF CONCEPTS (POC)



NOW THAT YOU'RE FAMILIAR WITH THE KEY CONSIDERATIONS AND STRATEGIES FOR TRANSITIONING FROM A RELATIONAL DATABASE TO A NOSQL DATABASE, YOU'RE READY TO START A PROOF OF CONCEPT. Now that you're familiar with the key considerations and strategies for transitioning from a relational database to a NoSQL database – how to select an application, how to model and access the data, and how to deploy the database – you're ready to start a proof of concept.

Couchbase solutions engineers have helped, and continue to help, many enterprises successfully introduce NoSQL, from planning all the way to post-production. We encourage everyone to start with a proof of concept.

There are five steps to a successful proof of concept:

1. Select a use case and application

It's important to remember that the key to successfully introducing a NoSQL database is to first identify an appropriate use case and select an application. Try to find an application that can realize one or more of the benefits of a NoSQL database – better performance and scalability, higher availability, greater agility, and/or improved operational management.

2. Define the success criteria

It may be difficult to move beyond a proof of concept without defining how to measure its success. Success criteria vary for different applications. For some, it may be performance (e.g., 5ms latency in the 95th percentile). For others, it may be management (e.g., easier to scale and add nodes). It may be faster development cycles. Whatever it may be, make sure to specify it upfront.

3. Understand the data

Before defining the data model, simply understand the data and the business domain. At first, the focus should not be on how to define or migrate the data model. Rather, it should be on understanding the data, independent of how it's stored in the database.

4. Identify the access patterns

Next, identify how the data is used and then begin to model it within a NoSQL database. This will depend very much on how the application reads, writes, and finds data. The data model can be optimized for different access patterns. In addition, you have to choose the right data access method – key-value operations, queries, full-text search, – for the right data access pattern – basic read/ write operations, queries, or aggregation and reporting.

5. Review the architecture

After completing the proof of concept and measuring the results against its predefined success criteria, it's time to begin preparing for production deployment. This is the time to review the architecture and create a blueprint for production. Based on the POC development experience, you can identify what worked well and what could work better – use this knowledge to help define the final application architecture. NoSQL was expressly designed for the requirements of modern web, mobile, and IoT applications. For enterprises that make the shift to NoSQL, the rewards are significant: greater agility, faster time to market, easier scalability, better performance and availability, and lower costs. Developers find that working with a JSON data model is far more natural than a rigidly defined relational schema, while operations engineers love the ease of elastically scaling the database without all the headaches of manual sharding and skyrocketing costs.

If you're ready to take the next steps and are looking for more specific advice, we invite you to talk with one of our solutions engineers. At a minimum, you'll probably get some helpful insights and best practices for your particular use case.

For enterprises that make the shift to NoSQL, the rewards are significant: greater agility, faster time to market, easier scalability, better performance and availability, and lower costs.





Modern customer experiences need a flexible database platform that can power applications spanning from cloud to edge and everything in between. Couchbase's mission is to simplify how developers and architects develop, deploy and consume modern applications wherever they are. We have reimagined the database with our fast, flexible and affordable cloud database platform Capella, allowing organizations to quickly build applications that deliver premium experiences to their customers—all with best-in-class price performance. More than 30% of the Fortune 100 trust Couchbase to power their modern applications.

For more information, visit www.couchbase.com and follow us on Twitter. © 2023 Couchbase. All rights reserved.

