

WHITEPAPER

How to Choose a Database for Your Mobile Apps



Contents

HOW TO CHOOSE A DATABASE FOR YOUR MOBILE APPS	3
Evaluating your mobile database – a checklist	3
SUPPORT FOR THE RIGHT PLATFORMS	4
SECURE AT REST AND IN MOTION	4
FLEXIBLE DATA MODELS	4
SYNC WITH THE RIGHT PARTITIONS	5
PESKY CONFLICTS	5
SYNC AT THE RIGHT TIMES	6
FLEXIBLE DEPLOYMENT	6
LOCAL DATA STORAGE CAPABILITIES	7
VECTOR SEARCH FOR AI FEATURES	7
SHOULD YOU BUILD OR SHOULD YOU BUY?	7
HOW THE MAJOR PLAYERS STACK UP	8
WHY COUCHBASE MOBILE?	8

HOW TO CHOOSE A DATABASE FOR YOUR MOBILE APPS

Evaluating your mobile database - a checklist

Successful mobile apps rely heavily on their database provider. How does your mobile database stack up when put to the test?

- Does it have support for the right platforms?
- Can you control your database sync?
- What sort of querying options does your local database support?
- Does it support vector search for semantic search and GenAl features?
- · Can you easily integrate with LLM or ML models?
- Is data secure at rest and in motion?
- Do you deploy on-prem or in the cloud?
- · Does the sync technology support multiple topologies?
- Do you worry about adapting to changing data schemas?
- How does your database handle pesky conflicts?
- Do you need to support multi-tenant apps?
- Should you build or should you buy?

The choice of database plays a key role in building successful mobile apps. Data synchronization, local data storage and querying capabilities, vector search for AI features, and end-to-end security are key elements of a data platform.

Today's consumers are highly reliant on their mobile applications, with large portions of modern business relying on the ability to interact with their users via mobile. If apps don't work, users won't use them – it's that simple.

To avoid reliance on a network, providers of databases and cloud services have started to add synchronization and offline capabilities to their mobile offerings. If apps require a connection to work, then the end-user experience will be sluggish and unpredictable.

Solutions like Couchbase Mobile, MongoDB[™] Atlas Device Sync (now deprecated to end-of-life), Amazon's AWS AppSync, and Google's Cloud Firestore offer the all-important sync that enables apps to work both online and offline. Other database providers such as SQLite, Android Room, and Core Data support local storage but do not provide sync support.

With so many offerings available, how does a mobile developer select the right technology? The following key criteria are essential when evaluating mobile solutions: multi-platform support, local data storage capabilities, sync capabilities with conflict resolution, ease of development, security, vector search for AI features, agile data modeling, flexible deployment, and topology options.



THE CHOICE OF DATABASE PLAYS A KEY ROLE IN BUILDING SUCCESSFUL MOBILE APPS. What client platforms are supported? Do you need to go beyond iOS and Android? Are you looking to support platforms that aren't traditionally considered mobile, such as embedded systems and IoT devices? Are you looking to support Windows and Mac desktops and laptops as well? What about cross-platform technology?

Many of today's applications start on mobile, then add a native desktop version.

It's important to evaluate database and cloud options based on the platform support that you need not only today but also in the future.

SECURE AT REST AND IN MOTION

When you're using synchronized and decentralized storage, it's important to access, transmit, and store data securely. To cover this completely, you need to address authentication, data at rest, data in motion, and read/write access control.

Authentication should be flexible and allow for the use of common public and custom authentication providers. Support for anonymous access is also important for many apps. For data at rest on the server and client, you'll want support for both file system encryption and data-level encryption. For data in motion, communication should be over a secure channel like SSL or TLS. For data read/write access, the database should include granular user and role-based access control (RBAC).

FLEXIBLE DATA MODELS

MANY OF TODAY'S APPLICATIONS START ON MOBILE, THEN ADD A NATIVE DESKTOP VERSION. Data modeling flexibility will dictate whether you can articulate the model requirements for your apps in an efficient and appropriate way. Today's mobile apps evolve at a very fast pace, and the flexibility of your model will dictate whether you can easily adapt as your requirements change in the future.

A new release of a mobile app with an updated data schema will require expensive database schema migrations to be performed on app launch, adding to your app startup costs. As an app developer, you don't have control over when a new version of your app gets adopted. As a result, users might be migrating from a very old version of schema to the latest version, exacerbating data migration issues.

Relational databases are still a good choice if an app requires strong data consistency or its data is highly relational. But NoSQL databases offer much greater flexibility. Configurable sync topology support is needed to allow you to meet your partition requirements. In other words, you need the ability to configure the system to allow certain parts to operate offline. A star network is the most common topology, where each device is connected to a central hub using a point-to-point connection that allows the devices to operate offline. Other common topologies such as tree and mesh allow different parts of the system (in addition to the devices) to operate offline.

You may also want support for cloudless topologies that allow devices to communicate peer to peer and directly sync data among themselves. Peer-to-peer synchronization is a powerful addition or alternative to client/server synchronization. It allows apps to connect and exchange data directly without going through a central server in the cloud. As a result, apps can continue to work and share data regardless of internet availability.

A point-of-sale (POS) system is a good example of a tree topology. POS systems require that a brick-and-mortar store continue to operate if it becomes disconnected from the rest of the system. In this configuration, POS devices would sync with a store-level database, which would then sync with a global system. So, stores can continue to operate and sync data with their POS devices regardless of connectivity to the global system.

PESKY CONFLICTS

For mobile platforms, or any other platform that utilizes decentralized data writes, the same data can be simultaneously modified on multiple devices, creating a conflict. The system needs to support a mechanism for resolving those conflicts. Conflict handling will differ for each system. Couchbase Mobile, for example, uses revision trees with a default resolution rule of "most active branch wins." This is the same approach taken by revision control systems such as Git and much different than clock-based systems that take a "most recent change wins" approach. Clock-based resolution systems are problematic due to the issues around clock differences across devices.



In addition to being able to resolve conflicts, it's important to have the ability to control how the system syncs, which includes replication strategy, replication events, conditional replication, and replication filtering. For replication strategy, look for support for streaming, polling, one time, continuous, and push. The granularity of sync has a direct impact on network bandwidth usage, so having a solution that is smart about identifying what subset of data has changed and only syncing the deltas is an important consideration in mobile deployments where data plans come at a premium and network bandwidth is limited.

You should also have the ability to use a combination of these strategies. For replication events, you may need to know the overall replication status as well as the replication status of individual documents. For conditional replication, you may need to replicate data only under certain conditions, such as when the device is on Wi-Fi or when it has sufficient battery power. For replication filtering, you should have the ability to replicate some data but not other data. The filtering could be fine grained and could be based on the content of the data itself.

FLEXIBLE DEPLOYMENT

AVOID VENDOR LOCK-IN SO THAT YOUR ABILITY TO GROW OR MIGRATE PROJECTS IN THE FUTURE IS NOT LIMITED BY FUNCTIONALITY OR COST ISSUES.



How you deploy the database is a very important decision and one that should not be constrained by requirements to use a particular hosting provider or platform. Avoid vendor lock-in so that your ability to grow or migrate projects in the future is not limited by functionality or cost issues. The ultimate solution should allow you to choose between using fully managed and hosted backend data and sync services, or hosting your backend database and sync yourself on any public or private cloud, as well as on premises in your own data center. Ideally, you would not need any additional third-party hosted services on additional software to build your complete mobile stack.

If your app needs to support multiple tenants – each with its own users – ensure that your database can isolate and secure data for each tenant at the appropriate levels of granularity, and without the need to create separate database instances.

Another way to ensure flexibility is to use modern containerization approaches that make it easy to manage deployments as needs change or grow. Using Kubernetes to help orchestrate your containerized environments makes it much easier for your operations team to keep things manageable.



When looking for offline storage capabilities, you'll need to determine if your app is expected to be usable in a completely standalone mode with extended periods of network disconnectivity, or if you're looking for a temporary cache to handle short network disruptions. An extensive database query API with SQL and full-text search support and the ability to be asynchronously notified of database changes will allow mobile apps to support responsive and highly reactive workflows locally. An intuitive, easy-to-use API will reduce ramp-up costs and reduce development and integration efforts.

VECTOR SEARCH FOR AI FEATURES

The search feature in your app should always return personalized responses in order to make an impact. But searching only for specific words and phrases is not enough to produce accurate results in context. To make a personal connection you need semantic search for matches that are meaningful to the user. Vector search finds related information based on the core meaning of the input, making it the best option for providing relevant information that connects with users.

And for GenAl functionality such as conversational chatbots, recommenders or Al-assistants, vector search enables easy integration with LLMs (large language models) through techniques such as Retrieval Augmented Generation (RAG) where current local vector data is passed along with prompts to provide better precision and context for LLM responses. To enable these kinds of Al features, make sure your database supports vector search, including within offline storage to meet strict privacy or regulatory requirements.

SHOULD YOU BUILD OR SHOULD YOU BUY?

When looking to add sync to your apps, you'll need to determine if you should build a solution or get it from a provider. Building sync correctly is notoriously difficult and expensive, as it must deal with all of the complexities of distributed computing. For most apps, you'll be better off leaving data synchronization to a specialized stack and focusing on your app features. The key is choosing a solution that is flexible. If you go down the build path, be ready to expend a significant portion of your time and resources on building sync and supporting everything listed above.

When choosing a mobile sync and storage provider, taking full measure of the above criteria will be critical to building secure, flexible, and manageable mobile apps that always work – with or without an internet connection.



HOW THE MAJOR PLAYERS STACK UP

Capability	Couchbase Mobile	MongoDB Atlas Device Sync	Google Cloud Firestore	AWS AppSync	SQLite, Core Data, Android Room
Offline support					
Platform support			-		
Enterprise-level security					_
Flexible data model					_
Flexible topology support			_	_	_
Peer-to-peer sync		_	_	_	_
Delta synchronization		_	_	_	_
Flexible deployment					_
Vector search on-device		_	_	_	_

WHY COUCHBASE MOBILE?

Couchbase Mobile brings the power and flexibility of a NoSQL database to the edge. It includes Couchbase Lite, an embedded NoSQL database for mobile and embedded apps that exposes a powerful SQL++ query API and supports vector search on-device as well as in the cloud. It also includes a synchronization gateway responsible for synchronizing data across clients and the cloud in order to enforce access control policies, authentication, authorization, and data routing. Choose from fully managed and hosted sync with Capella App Services, or host and manage sync yourself with Sync Gateway. This powerful combination, especially when paired with Couchbase Capella[™] DBaaS, is what makes Couchbase Mobile the only choice for secure, resilient, offline-first applications that deliver sub-second responsiveness and 100% uptime.

Learn more at www.couchbase.com/mobile, and sign up for the Capella free trial at cloud.couchbase.com/sign-up.





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 X (formerly Twitter) @couchbase.

© 2024 Couchbase. All rights reserved.

