SAMPLE RFP: NoSQL DATABASE EVALUATION CRITERIA

Couchbase

There are many different criteria that go into selecting a database for your application. From development through to performance, scale, and management, you need to choose a database that matches the requirements needed to unleash the power of your application. If you've determined you need a NoSQL database for your application, the criteria below will explain how you can evaluate Couchbase and other N0SQL alternatives based on your requirements.

Evaluation Criteria	
1	Data model will evolve over time and/or data will need to be aggregated from multiple sources
2	No downtime during application releases, including updates to data model
3	Application needs to access data in multiple ways (i.e., structured query, free-form text search, graph traversal)
4	Application developed in Node.js, Go, PHP, Python, Java (with or without Spring), or .NET (with or without Linq)
5	Application requires SQL and/or joins on JSON
6	Application requires Kafka, Spark, or NiFi
7	Application requires mobile capabilities including local database, offline data access, and sync
8	Application performance and consistency required during spikes in usage
9	Low-latency and/or high-throughput access
10	Operational analytics
11	One single technology for different workload types (query, text search, etc.)
12	Deployed or will deploy on any or all of the public clouds
13	Deployed or will deploy within containers (Docker, Kubernetes, OpenShift)
14	High availability (HA) to meet stringent uptime SLAs
15	Geographic data replication for disaster recovery (DR) and/or for pushing data closer to users at the edge
16	100% uptime for planned maintenance such as patching or upgrades
17	Eliminate ETL for faster data insights
18	Security (including GDPR, HIPAA, FIPS) of user-facing data

CRITERIA 1:

Are you building a new application with the expectation that the data model will evolve over time? Or do you need to aggregate data from multiple sources?

COUCHBASE SOLUTION:

Couchbase is a document database that requires no schema definition to store data. Couchbase documents are JSON, a self-describing format capable of representing rich structures and relationships. Unlike a traditional RDBMS, the schema in Couchbase Server is a logical construct entirely defined in the application code and captured in the structure of the stored documents. Since there is no explicit schema to maintain, developers can add new objects and properties at any time just by pushing new application code that stores new JSON, without having to also make the same changes to the schemas. This allows applications to evolve quickly and seamlessly.

Further, the flexible schema model makes it easy to aggregate data from multiple data sources into a single service layer built on Couchbase.

CRITERIA 2:

Is your application suffering from performance issues due to ORM or other inefficient ways of converting relational structures to/from JSON? Or is your application development/ deployment schedule impacted by time-consuming data model reviews? Is your ability to release new versions of your application impacted by downtime requirements to update the data model within your database?

COUCHBASE SOLUTION:

Couchbase stores data as JSON, which is flexible and does not require upfront schema definitions, and also has the power to capture structure and relationships as required by the application. Couchbase's query language, N1QL, can impose schema-on-read to provide relational SQL-like processing using the structure of the JSON. This combination of JSON and N1QL enables agility in application development while retaining the power to capture and process relationships in the data.

In addition, Couchbase also provides an INFER command via N1QL that helps understand the schema of JSON documents, even in a large dataset, by sampling.

CRITERIA 3:

Does your application need to access data in multiple ways (i.e., structured query, free-form text search, graph traversal)?

COUCHBASE SOLUTION:

Couchbase offers rich and comprehensive support for a broad range of query types:

- Structured relational SQL-like queries using Global Secondary Indexes (GSI)
- Unstructured text queries using Full-Text Search (FTS)
- Ad hoc queries over large amounts of data for real-time insights using Couchbase Analytics

Couchbase N1QL uses Global Secondary Indexes and supports most SQL functionality such as joins, subqueries, GroupBy as well as JSON-specific features such as nested path expressions, NEST, and UNNEST. N1QL is actually the first commercial implementation of SQL++ which encompasses both SQL and the JSON data model.

The following is a link to an ebook on N1QL: http://blog.couchbase.com/wp-content/uploads/2017/03/N1QL-A-Practical-Guide-v2.pdf.

Couchbase Full-Text Search includes indexing, faceting, scoring, and text analysis for 20+ languages. FTS provides 20+ query types in 7 categories: simple, compound, range, query string, geospatial, non-analytic, and other special queries.

Couchbase Analytics supports completely ad hoc queries using N1QL (including full scans) for operational insights.

CRITERIA 4:

Is your application developed in Node.js, Go, PHP, Python, Java (with or without Spring), or .NET (with or without Linq)?

COUCHBASE SOLUTION:

Couchbase offers numerous software development kits (SDKs) that enable programmatic access to Couchbase Server. The SDKs include:

- C SDK
- Java SDK
- .NET SDK
- Node.js SDK
- Go SDK
- PHP SDK
- Python SDK

CRITERIA 5:

Are your developers familiar with SQL? Do you need joins on JSON?

COUCHBASE SOLUTION:

Couchbase N1QL is a query language based on SQL that is optimized and extended for JSON. N1QL provides rich join functionality that allows combining data with multiple relationships. N1QL supports all the classic SQL joins including Inner Join, Left Outer Join, Right Outer Join, and subquery with joins in from clause. Refer to this DZone article to learn more.

Couchbase 5.5 also added support for full ANSI joins on arbitrary fields that are not necessarily parent-child relationships. The following is a detailed blog on the ANSI join support: https://blog.couchbase.com/ansi-join-support-n1ql/.

CRITERIA 6:

Are you receiving data from and/or sending data to Kafka, Spark, or NiFi?

COUCHBASE SOLUTION:

Couchbase connectors enable applications to exchange data with a number of other platforms including modern big data platforms such as Kafka, Spark, and NiFi.

CRITERIA 7:

Are you developing a mobile or embedded application? Do you need to provide offline access to data?

COUCHBASE SOLUTION:

Couchbase fits well as a scalable online platform into the ecosystem of these big data technologies by providing source and sink connectors to them.

Couchbase Mobile seamlessly extends the Couchbase to the edge, securely managing and syncing data from any cloud to every mobile device. Couchbase Mobile is the full-stack combination of Couchbase Lite, the enterprise-grade embedded database on device, Couchbase Sync Gateway, the middle-tier replication layer, and Couchbase Server, the enterprise-class NoSQL database. In addition, Couchbase Lite can operate offline as a standalone embedded database, replicating directly between devices if needed. Learn more about Couchbase Mobile at: https://www.couchbase.com/products/mobile.

CRITERIA 8:

Are you unsure what the adoption or user growth of your application will be? Does your application performance degrade during peak times? Is inconsistent performance of your application during peak times causing customers to go elsewhere?

COUCHBASE SOLUTION:

Couchbase is a scale-out, globally distributed system with a simple and flexible architecture to allow scaling in multiple ways to meet an application's evolving needs.

- Couchbase automatically shards data across nodes. New nodes can be added online (without any downtime) and data will automatically be rebalanced across all the nodes in the cluster. This elasticity allows for capacity to be easily scaled (up or down) either for long-term growth or to account for temporary peaks.
- Couchbase allows independent scaling of each service including Data Service, Index and Query Service, Full-Text Search and Analytics. This enables more optimized scaling targeted only for workloads that require it.
- Geographic distribution using XDCR (cross datacenter replication) enables global user growth by bringing data closer to users across the globe.

CRITERIA 9:

Do you require low-latency and/or high-throughput access?
Or are you looking for an in-memory database or a caching layer?

COUCHBASE SOLUTION:

Couchbase, with its memory-first architecture and tightly managed caches, provides sub-millisecond read latencies with high predictability. With asynchronous durability, Couchbase provides sub-millisecond write latencies as well. This translates to throughput exceeding a million operations per second and per node on commodity servers.

In addition, Global Secondary Indexes allow indexes to be partitioned independently of data, hence minimizing latency for multiple access patterns.

Benchmarks of Couchbase read/write performance in comparison to MongoDB™ and DataStax are available at https://www.couchbase.com/benchmarks

---

CRITERIA 10:

Do you need to perform operational analytics?

COUCHBASE SOLUTION:

Couchbase Analytics provides real-time operational insights by constantly ingesting data in real time from the Couchbase data nodes. It has a massively parallel processing (MPP) query engine that is capable of executing queries in parallel on multiple servers and leveraging all the cores on each server as needed to process large amounts of data efficiently.

Couchbase Analytics also provides BI visualizations via partnerships with JSON-based tools such as Knowi and Power BI as well as more traditional ODBC/JDBC integration using cdata drivers for tools like Tableau, Looker, and MicroStrategy.

CRITERIA 11:

Are you managing multiple products for different workload types (query, text search, etc.)?

COUCHBASE SOLUTION:

Couchbase not only provides a high speed and scalable document store, but also provides integrated functionality for query, text search, geospatial search, and analytics. Additional features include:

- Built-in query language N1QL for scanning, filtering, sorting, aggregating, and joining JSON data using familiar SQL syntax.
- Full-Text Search (FTS) engine to support text search including faceted navigation on data that is stored in Couchbase.
- Geospatial queries and creation of geospatial indexes. Geospatial querying allows nearest neighbor queries, distance/radius-based querying, and bounding box-based querying.
- Analytics for ad hoc queries over large amounts of data.

Each of these services can be scaled independently, but operated and monitored in an integrated manner from one console.

CRITERIA 12:

Are you deploying within a public cloud? Or are you considering moving to a public cloud? Do you need the flexibility to move between public clouds now or in the future?

COUCHBASE SOLUTION:

Couchbase's cloud strategy is about the choice and flexibility of deployment, while giving customers complete control and ownership of their data. Couchbase was built using a cloud-native architecture as a scale-out distributed database that can be deployed anywhere across physical, virtualized, public cloud or containerized infrastructure, as well as Couchbase managed services environments. Couchbase's cloud offerings leverage these architectural advantages to focus on enabling customers to address their hybrid and multi-cloud requirements with a consistent data platform that can be deployed across on-premises, edge and cloud deployments.

In addition, Cross Datacenter Replication (XDCR) can be used to replicate data between deployed clusters in the cloud.

CRITERIA 13:

Are you deploying (or considering to deploy) within containers (Docker, Kubernetes, OpenShift)?

COUCHBASE SOLUTION:

Couchbase supports Docker containers and orchestration using Kubernetes for easy deployment and management. Couchbase is the first NoSQL database to offer production support for Docker (version 4.5), and the first NoSQL database to offer an Autonomous Operator for Kubernetes and OpenShift (version 5.5).

CRITERIA 14:

Do you need high availability (HA) to meet your stringent uptime SLAs?

COUCHBASE SOLUTION:

Couchbase high-availability architecture with automatic failover ensures always-on access to your data. Features include:

- Automatic balanced setup and replication of data to a configurable number of replicas
- Fast and reliable automatic failover to ensure minimum downtime to application
- Built-in rack awareness for protection from co-related node failures
- Cross datacenter replication (XDCR) to guard against complete datacenter failure

---

CRITERIA 15:

Do you need to replicate data geographically for disaster recovery (DR) or for pushing data closer to users on the edge?

COUCHBASE SOLUTION:

Couchbase XDCR allows replication to multiple, geographically diverse datacenters either for disaster recovery or to bring data closer to users (e.g. mobile devices, small clusters on cruise ships, etc.) for faster data access.

XDCR can be set up in an active-active configuration and allows multiple strategies for automatic conflict resolution.

XDCR supports many multi-master topologies including star, ring, chain, mesh, one-to-many, and many-to-one.

CRITERIA 16:

Do you require 100% uptime for planned maintenance such as patching or upgrades?

COUCHBASE SOLUTION:

Couchbase supports online upgrade without any downtime to the application. Depending on whether spare nodes are available, the online upgrade can be done either as a full swap rebalance, a rolling swap rebalance, or as a traditional rolling upgrade. Swap rebalances are the recommended approach because they are 100% online, fast, and have minimal impact to performance of the application. The following is a helpful blog on rolling upgrades.

CRITERIA 17:

Are you using ETL from one or more databases for operational data access?

COUCHBASE SOLUTION:

Couchbase Analytics requires no ETL from the operational data store to perform analytics. In the Couchbase Data Platform, data is automatically streamed to the analytics data store in real time, hence enabling real-time insights without the latency and overhead of ETL.

CRITERIA 18:

Are you concerned with security (including GDPR, HIPAA, FIPS) of your user-facing data?

COUCHBASE SOLUTION:

Couchbase provides secure data everywhere – on the wire, on the device, in the cloud, and in the datacenter and covers all aspects of securing customer data including authentication, authorization, auditing, and encryption. Authentication allows multiple methods including certificate-based, password-based, LDAP, and PAM-based. Authorization is fine-grained with Role-Based Access Control (RBAC). Couchbase provides native encryption on the wire and third-party integration for encryption at rest.

A complete list of Couchbase security best practices can be found here: https://www.couchbase.com/resources/security.

About Couchbase

Couchbase's mission is to be the database platform that enables a revolution in application innovation. To make this possible, Couchbase created an enterprise-class NoSQL database to help deliver ever-richer and ever more personalized customer and employee experiences. Built with the most powerful NoSQL technology, Couchbase was architected on top of an open source foundation for the massively interactive enterprise. Our geo-distributed database provides unmatched developer agility and manageability, as well as unparalleled performance at any scale, from any cloud to the edge.

Couchbase has become pervasive in our everyday lives; our customers include industry leaders Amadeus, AT&T, BD (Becton, Dickinson and Company), Carrefour, Cisco, Comcast, Disney, DreamWorks Animation, eBay, Marriott, Neiman Marcus, Tesco, Tommy Hilfiger, United, Verizon, Wells Fargo, as well as hundreds of other household names. For more information, visit www.couchbase.com.

3250 Olcott Street | Santa Clara, California 95054 | 1-650-417-7500 | www.couchbase.com