



**THE TCO ADVANTAGES OF COUCHBASE
CLOUD'S IN-VPC DEPLOYMENT**

Executive Summary

[In January 2020, Forrester Research vice president Jeffrey Hammonds observed](#) that the reason enterprises have only moved twenty percent of their data-intensive applications to the cloud is due to how the stewards of that data in IT treat their information assets as something precious, like a family pet rather than viewing it as a life-sustaining, but replenishable resource like food. In order to help IT make this shift in mindset, we must be able to ensure them that they will not sacrifice existing controls over their data assets by shifting them to the cloud. And, due to economic events that occurred three months later, we must recognize that enterprises of all sizes are optimizing their decisions towards total cost of ownership because containing costs has returned as the overriding factor in IT decision making.

This paper describes how organizations can minimize and control cloud operating costs of NoSQL database deployments. It is intended for developers, cloud architects and DevOps professionals tasked with deploying or migrating applications and underlying database software into cloud infrastructure in a cost effective manner.

This document will construct a cost savings framework that can be used to estimate and evaluate database service providers' total cost of ownership advantages over one another. It will offer a variety of dimensions or areas of focus using Couchbase Cloud as its example. Where appropriate, it will identify questions to pose to alternative vendors competing for the customers' business.

The Hypothesis

In-VPC deployment is an emerging best practice technique used by independent software vendors (ISV) such as Databricks, Cloudera and Couchbase to offer customers the nearest proximity of resource, asset and cost control to what they enjoyed when managing their own custom-built data centers. To this premise we will combine unique capabilities of Couchbase Cloud, including its performance results in head-to-head comparisons, and then demonstrate that Couchbase, Inc. has developed a transparent and predictable framework that can be used by customers to evaluate the cost-effectiveness of competing solutions in delivering the lowest possible TCO.

The Method

The paper will break down the functional and cost control benefit of this In-VPC deployment technique using a number of dimensions that effect a customer's ability to control their costs of operation. These dimensions fall into the following categories:

- **In-VPC Deployment Strategy:** Will you combine the database and infrastructure services or not? We offer multiple reasons why choosing an ISV such as Couchbase that hosts their managed service in the virtual private cloud of your choice is the best alternative for cost savings and risk reduction. We will also explain why choosing a cloud service provider's built-in database only works to favor the CSP, and not you as the customer.
- **Couchbase Performance and Scaling Capabilities:** Is your database up to task of being efficient and maximizing the resources it has available? This category will examine cost saving gains due to raw node-for-node performance as demonstrated through benchmarks and due to elastic cluster rebalancing, scaling features such as cross data center replication (XDCR) and multidimensional scaling (MDS) of its independent services. This will allow you to performance-match your application to your database, and that to its infrastructure, ensuring that you can maximize your infrastructure resource consumption.
- **Licensing and Transparency Policies:** Does your vendor have your best interests at heart? This section will detail specific policies that Couchbase has enacted that effect your cost of ownership versus other providers. Transparency in this regard is turned into dollars for the customer.

We will next decompose each of these dimensions to examine the tradeoff choices available to the customer that effect cost control. Finally, we will provide you a tabulation framework from which you can evaluate your price performance and ownership costs benefits of Couchbase Cloud versus its competitors.



With Couchbase Cloud, the customer remains in control of their data, it's security, it's network and the Couchbase clusters hosting it at all times.

The In-VPC Deployment Strategy

In-VPC deployment is the practice of deploying fully-managed software applications, usually containerized and managed by Kubernetes-based tools, into the customer's own virtual private clouds. It is an advantage typically held by cloud service providers (CSPs) who also offer software products in addition to hardware and networking infrastructures as a service (IaaS), but today, select ISVs are adopting this deployment strategy as they make packaging, installation, management and upgrades easier through technologies such as containers and Kubernetes. Couchbase believes that this is an emerging best practice that will accelerate to become a new cloud norm because it provides advantages last seen when enterprises ran data centers of their own.

The In-VPC Deployment Strategy is designed to offer customers higher degrees of control over resources, data security and configurations than other alternatives such as choosing to run a traditional Software-as-a-Service (SaaS) application that blends together its infrastructure and the vendors' software as is seen by competing Database-as-a-Service (DBaaS) providers such as Neo4j and MongoDB, who bundle your hardware and software delivery together, as do CSPs who offer products such as Microsoft's CosmosDB, Google's Bigtable or DynamoDB from Amazon Web Services (AWS).

Couchbase believes that the decision to single-source your database from one vendor versus keeping IaaS and database separate appears tempting, but it should be reconsidered. Your first concern is data ownership and access. Your second concern is infrastructure choice, configuration and management, and your third concern is your primary cost containment driver, whose negotiating strength do you trust for determining your infrastructure costs? And finally, what kind of safety net do you have if you seek to change cloud providers? In-VPC deployment provides satisfying answers to each of these questions.

Data Security and Sovereignty

Are you the exclusive owner of your cloud data? Do you and your IT staff determine how that data is secured, behind what firewalls and using which encryption protocols? When you choose a vendor who bundles infrastructure with your database software, you are giving up control over how and where that data is secured to that vendor. With In-VPC deployment, that control remains in your grip.

Further, if you operate a global business, you are subject to a variety of data sovereignty (where and how data physically resides) regulations such as GDPR or even CCPA if you store personally identifiable information. With In-VPC deployment, you choose where the data resides by selecting VPC accounts in the local regions or zones offered by your CSP. Couchbase Cloud customers can choose from multiple regions, and by 2020's end, from multiple cloud service providers, where to install their clusters. This eases the burdens of adhering to local regulations. By contrast, data sovereignty may be limited by vendors who bundle IaaS within their offerings.

With Couchbase Cloud, the customer remains in control of their data, it's security, it's network and the Couchbase clusters hosting it at all times. This design creates a number of secondary benefits, including:

- The ability to secure and isolate their data as they wish, and to "lock out" third party access, including Couchbase, if necessary.
 - Isolation could be scoped on a per-cluster, per tenant or other means.
 - Also allows for private subnet definition for backend services when applications are deployed to the public internet.
 - Implement their preferred network security layering strategy, including using security groups, network access control lists, TLS and other IT-enforced policies and techniques to further secure access to Couchbase resources.
- The ability to specify IP address ranges, subnets, routing tables, and network gateways according to their own IT policies.
- Easing compliance burdens to operational control standards like SOC Type II and to data privacy and sharing regulations like HIPAA & CCPA without requiring special exceptions like a Business Associates Agreement letter (BAA Letter).



Couchbase offers over a dozen supported instance profiles, which allows customers to know exactly the instance type and its associated operating costs.



Hiding Infrastructure Choices and Configurations

Some vendors obfuscate the infrastructure and instances offered to the customers without allowing them to know exactly for what they are paying. This is very common among early generations of SaaS providers whose package includes the hosting and operation of the infrastructure running their managed software service. While these organizations may argue that this affords them the ability to tune the system down to bare metal, it also means that they must pass on to the customer the capital and operational expenses of hosting a data center or tuning a public-cloud provider's environment to their specifications. These costs are almost certainly hidden from the customer in order to avoid the uncomfortable conversation about whether the SaaS provider is building in a profit margin to their infrastructure (a packaging gimmick called, "margin stacking" the stack). An equally difficult conversation is settling on whose purchasing department can negotiate with the CSP more effectively, the customer or the vendor?

Couchbase is transparent in allowing the customer to choose the infrastructure upon which to run their cluster. Couchbase offers over a dozen supported instance profiles, which allows customers to know exactly the instance type and its associated operating costs. Knowing instance types eliminates the opportunity for margin stacking, and Couchbase operations can be added as a private offer to the customers' CSP monthly bill if they desire.

Negotiate the Best Price with Reserved Instances

Databases are often systems that run without downtime, which opens the opportunity for the customer to save on infrastructure costs by licensing reserved instances, annual subscription contracts for infrastructure offered by CSPs, which can dramatically lower their ownership costs over on-demand consumption. In fact, [AWS advertises savings of over 75%](#) for EC2 instances.

It is also likely that the enterprise runs a significant number of instances with their cloud provider, and has negotiated special pricing for those instances, so it is advantageous for the customer to be able to apply their organization's discounts for the instances they allocate for Couchbase Cloud. This is an advantage that is lost when vendors choose the infrastructure for the customer.

Choose a Database ISV with Multicloud Support

An additional advantage of choosing an ISV versus products directly offered by cloud service providers is the ability to support multiple clouds, which helps customers avoid being over-leveraged by a single cloud vendor. Multicloud support, provided by ISVs, helps enterprises to avoid the long term risks of dependence on a single cloud provider. This is why organizations should consider if choosing CosmosDB locks them to Azure, DynamoDB locks them to AWS, or Bigtable to Google Cloud Platform (GCP)?

Couchbase Cloud can act as a safety hatch for organizations who need to change from one CSP to another for any number of reasons. Couchbase will support AWS, Azure and GCP by the 2020 year-end. And further, Couchbase Server and its Kubernetes Autonomous Operator allows you to deploy and manage Couchbase Clusters in custom-constructed virtual clouds, Oracle Cloud Infrastructure, VMWare's Pivotal Cloud Foundry, IBM's Red Hat Openshift and so on.



Node for node, Couchbase is the performance leader among popular NoSQL databases with throughput speeds 2 to 3 times faster than competitors.

Couchbase Performance and Scaling Capabilities

Costs of operations is directly correlated to the efficiency of software. Some vendors are highly motivated to maximize performance for competitive gains while some are less motivated to focus on the arduous process of performance testing and development. This section will look at the effects of specific features of Couchbase Cloud and describe their effects on operational costs. It will explain labor cost advantages achieved in using the Couchbase Control Plane as well as why software workload performance, design and tuning techniques, replication, scaling and rebalancing features all contribute to lowering operational costs.

Single Pane Interface for Managing Clusters and Clouds

DevOps managers need comprehensive oversight tools to help them keep clusters and clouds running smoothly. The challenge to overcome here is to avoid increasing management complexity at the same rate as a deployment grows. Instance and cluster management must allow one DevOps administrator to oversee large scale, distributed clusters and instances all at once. The cluster control plane should be omniscient not only among clusters deployed regionally, but also be aware of connected clusters in other zones or even other clouds.

The Control Plane should also support all the required maintenance of the cluster including managing installation, configuration, backups, fail-over and upgrades of any cluster member. It should oversee the automatic, elastic scaling and data rebalancing capabilities as cluster nodes are added or decommissioned. This allows DevOps managers to help match availability to demand during peak load and return to normal afterward. The control plane should also provide operational health information to the DevOps manager and export log information to external systems management dashboards such as Prometheus. All of this improves the effectiveness of the DevOps administrator in managing clusters such as Couchbase.

Software Workload Performance

Some vendor's software can run the desired workload faster or more efficiently than its competitors given a uniform operating platform. This is often demonstrated via head-to-head performance benchmarks to demonstrate the assertion. For example, in 2018 Altoros, a software implementation services company, [published workload benchmark between Couchbase Server, MongoDB and DataStax](#). Couchbase was able to achieve operation throughput speeds that were 2 to 3 times faster than these competitors across a variety of tests. These benchmarks demonstrated that Couchbase is, node-for-node the performance leader among popular NoSQL databases. Imagine the savings if your workload could run at an acceptable level with lightly provisioned hardware, dropping CPU cores from 8 to 4 or lowering memory requirements which could lead to 40 - 50% additional savings per cluster node.

What is interesting here is that in some cases, a SaaS vendor who also supplies infrastructure, especially those with a market-leading position, may not be overly motivated to optimize the performance of their service. Hints of this possibility can be found in their licensing policies (discussed below) such as charging hourly rates for their SaaS plus infrastructure service. In this case, they are able to make more money by allowing mediocre performance due to both slow software and aging infrastructure. A better solution is to decouple the software service from the infrastructure service and allow the customer to choose their performance profile and control the expenses allocated to each, just as they did managing a datacenter.

Performance-Oriented Database Features

Performance architectures must be devised at the outset of database design because it can be impossible to change the foundation of the database once it is established. For example, the Couchbase design is a byproduct of two complementary database architectures; an exceptionally fast memory-first design that performs operations in RAM in milliseconds, with the flexibility and scale of a JSON-structured document database. This combination sets Couchbase apart from databases that support one or the other of these designs. Today, for example it is not uncommon to use a memory-first key-value database as the caching system for a lesser performing document database that does not support query caching. The need to combine these capabilities with separate systems creates unnecessary degrees of complexity for development teams.



We offer Multidimensional Scaling (MDS) which is the ability to performance-match multiple service operations, such as indexing, querying, data control across one or many nodes in the cluster such that they are each able to meet the demands of the application. This is not offered in most other competitors.

Further, some document database vendors add other non-native functionality such as full-text search or analytic engines or mobile SDKs that can slow down development projects while also increasing the QA and debugging complexity of maintaining these types of applications. Clearly these vendors are taking architectural and feature cues from the Couchbase multi-service design, which not only includes native services for indexing, query and data control, but also adds multi-parallel processing analytic services, full-text search, event streaming and robust peer-to-peer synching mobile services. The combination of which is unmatched by any vendor, and the inspiration for the Couchbase NoEQUAL slogan.



Per Node Tuning Techniques

Some vendors do not offer the ability to tune how much or how few of its service resources run on any given node instance. Microservices-based deployments are gaining in popularity for a number of reasons, but primarily the desire to iterate rapidly through continuous integration and continuous deployment (CI/CD) practices is the driver. But there is also a cost control driver that can be found in a multi services-based approach to software development, and that is the ability to deploy services in a manner that matches the performance requirements of the application.

The opposite of a services-base design approach is to design an all-encompassing monolith of an application and simply scale it by deploying new nodes. PostgreSQL is an example of this as is MySQL and traditional relational databases as well as many NoSQL solutions. The issue here is that their one-size-fits-all configuration, multiplied by node, does not offer the ability to tailor the deployment based on the application's real workload. For example, if the application is excessively stressing the concurrency limits of the software, then invest in networking and connections. If the application is regularly repeating queries, and needs them quickly, then invest in memory and query result caching, or if a high volume of queries for specific result values are in play, then boost the database's indexing capacity. The key is to offer enough granularity to allow for core database services to be configured and assigned hardware resources based on the application's utilization of each service. Couchbase calls this "performance-matching" the database to the application.

Analytic vendors may argue that they do offer this services-based approach by articulating the value of separating their aggregation and computing resources from their data storage and retrieval processes, (good idea), and they would be correct to a degree, which is still beneficial to the customer in their ability to match their performance requirements to their infrastructure configuration. The Couchbase solution is to offer the ability to performance-match multiple service operations, such as indexing, querying, data control across one or many nodes in the cluster such that they are each able to meet the demands of the application. In Couchbase, this capability is called Multidimensional Scaling (MDS) of its services as a fine-grained performance control not found in most other competitors.



At Couchbase we offer discounted, pre-paid, annual consumption credits with no monthly minimum.

On-Prem to Cloud Replication, Sharding and Rebalancing

Couchbase also helps reduce costs by supporting a number of data replication and partitioning features. These include Cross Datacenter Replication (XDCR) which helps reduce network latency by placing data closer to the application user. It can also be used to filter data based on location in order to support residency requirements mentioned above. But its uses do not stop there. XDCR can drive backups to the cloud of on-premises deployments, or facilitate migrations into the cloud, or even culling live data sets for testing algorithm accuracy to mirror real-world conditions.

Couchbase also automatically supports sharding or partitioning of its database into logical buckets that can be easily moved or returned from one node to another in support of its elastic scaling and contracting based on the customer's workload. Subsequently, this auto-rebalance feature allows customers to identify the optimal cluster/node/service/instance configuration needed for their application and rebalance their existing workloads into it without disrupting their operation.

Transparent Policies for Purchasing, Usage and Licensing

Margin stacking may not be the only trick up a vendor's sleeve for peeling profits for themselves without being caught by the customer. Some vendors impose restrictive policies regarding levels of support; monthly minimum consumption; distinction of use, such as development and testing; cloning, co-locating or operating redundant systems for availability and backup, and more. Recognizing and avoiding vendors who adopt customer unfriendly policies should be avoided. This section discusses the cost advantages of specific policies of Couchbase.

Purchasing policies

Some vendors may margin stack by charging higher rates on the infrastructure they provide over what the customer might negotiate, which is why decoupling infrastructure licensing from software is so important. Another common practice, originally made famous by cellular service providers, is to require minimum monthly consumption levels in order to maintain their own revenue predictability. Here, they may set up their billing such that a customer is required to consume a minimum amount of their service per month, and if they are unable to meet that minimum, then they forfeit the unused resource. Similarly, if they go over the minimum usage limit they may be subject to overage charges. Couchbase counters this practice by offering discounted, pre-paid, annual consumption credits with no monthly minimum, which encourages the performance matching to the infrastructure and encourages the customer to grow into their deployments without making large resourcing commitments early in the project. The discount offered on these annual credits starts at 20%.

Usage Policies

Couchbase offers a number of customer-friendly policies for Couchbase Cloud that should be evaluated against competitors. For example, Couchbase allows the customer to choose service levels on a cluster by cluster basis, allowing them to select between 10x5 with 12 hour guaranteed response time for the Developer Pro package and for 24x7 with 1 hour response time for the Enterprise package, which is appropriate for mission-critical deployments.

This translates into a deeper level of control not likely found with other vendors. The ability to select your service levels on a cluster-by-cluster basis allows you to determine for which projects you need full-service enterprise support, and which ones do not.

Couchbase Cloud's Cost Savings Framework

The following formula building framework, will help organizations derive the cost savings advantages of Couchbase Cloud over its competitors. The framework consists of variables from topics explained earlier.



Category	Couchbase Cloud	MongoDB Atlas	Amazon DynamoDB
In-VPC Deployment Strategy			
In-VPC Deployment	4	0	4
Unbundled IaaS & Software (anti-margin stacking & optimization hiding)	4	1	1
Customer-Controlled Security	4	2	4
Choose your instance	4	3	4
Choose your CSP	3	4	1
Choose your region	3	3	4
Data Sovereignty & Regulatory Support	3	2	4
Choose your Cloud	4	4	1
Leverage your Negotiating Power	4	0	4
Performance and Scale of Couchbase			
Control all clusters from single console	4	3	3
On-Prem replication to and from clouds	4	4	2
Workload Performance node-for-node	4	2	2
Services-Based vs. Monolithic Design	4	1	1
Multidimensional scaling of services	4	0	0
Elastic node scaling & rebalancing	4	3	2
Transparent Policies for Purchasing and Usage			
No Minimum Monthly Consumption Fees	4	1	4
SLA-based Licensing per Cluster	4	2	2
Totals	65	35	43

In using this table, feel free to apply specific percentage-based weightings or specific costs and cost percentages to each category to factor in your own situation.



Conclusion

This paper has presented a compelling list of considerations when evaluating factors that contribute to Total Cost of Ownership for fully-managed databases as a service. Unlike most TCO or ROI frameworks that consider costs of labor saved within their argument, this one does not include such considerations. This framework evaluates the design and deployment decisions made by the vendor, the performance capabilities of their software, the scaling and tuning features within their software and the policies they apply to that software. And while it uses Couchbase-specific criteria, it can easily be modified to add criteria presented by other vendors in an attempt to create a balanced argument. Customers should use this framework to inform themselves of all the factors that contribute to the operating costs of fully-managed databases, and choose the vendor who best serves their interests. This paper will be updated regularly based on market condition, feedback and gathered intelligence.

For more information contact [Couchbase.com](https://couchbase.com) and start your free trial today.

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.