S&P Global
Market Intelligence

Couchbase expands Capella from managed database service to full-fledged data and AI platform

Analysts - James Curtis

Publication date: Tuesday, January 28 2025

Introduction

Couchbase Inc. has put significant development effort into its managed cloud service, Couchbase Capella. These development efforts push Capella well beyond a mere managed database service and into full-fledged database platform territory, catering not only to application development but also business analysts. Some current and near-recent updates include native analytics support for operational JSON data, updates to Couchbase Mobile with vector search for running AI-powered applications, a free environment tier for developers to test new features, and a full suite of AI services addressing data preparation, embedding models, integrating large language models and managing AI agents.

The Take

When Couchbase goes about its database development efforts, it has historically done so in the right way, even if it requires delaying some functionality until it is ready. A case in point is the rollout of analytics support on operational JSON data in August 2024. Couchbase added native columnar tables, which is the architecture used for traditional business intelligence analytics. This makes this effort all the more notable, given that Couchbase is a nonrelational operational database at its core. The company is not yet ready to take on the data warehouse titans just yet, but the foundation is there, given that integrated analytics support on operational JSON data is an underserved market. The company's efforts to add embedded AI services follow a notable trend many database vendors are pursuing and further give relevance to Couchbase as it continues to pursue the developer crowd that is pulling in AI for modern application development. As a whole, Couchbase is viewed as a rising player now, with its Capella platform targeting the operational and analytics database space. With continued development innovation, that trajectory can continue.

Context

Couchbase is an established player in the broader nonrelational, operational database market. Not quite four years out from its IPO in July 2021, Couchbase reports respectable numbers. From its third-quarter 2025 results, Couchbase recorded annual recurring revenue of $220 million, up 17% from in 2024. The company reports over 900 paying customers, of which one-third are part of the Fortune 100, suggesting that Couchbase continues to make inroads into larger enterprises. Couchbase Capella, its managed database cloud service, is now adopted by more that a third of Couchbase customers, revealing that the company still has notable growth given its on-premises customer numbers.

Couchbase looks well positioned to take advantage of the expected trajectory of NoSQL — the fastest-growing segment among the database areas that 451 Research tracks. Total forecasted revenue for the NoSQL market is expected to exceed $17 billion with an annual compound annual growth rate of 18.5% through 2028, according to the latest forecast from 451 Research's Data Platforms Market Monitor & Forecast 2024.

Technology updates

Couchbase rolled out columnar analytics support of JSON data for Capella in August 2024. The approach that Couchbase took is noteworthy, given that Couchbase Capella is a NoSQL database that is primarily geared toward operational applications. Most analytic databases, particularly data warehousing systems, deal primarily with highly structured relational data. JSON data is considered semi-structured and getting it in a structured format can be challenging when traditional, columnar-based systems expect structured data. As the name suggests, Couchbase implemented a columnar storage engine with a separate massively parallel processing engine with a cost-based optimizer. As JSON or relational data comes in — which can be ingested or streamed from a variety of sources, including Kafka — Capella quickly parses the data by recognizing its data type and stores it in columnar storage for immediate access. Users can then use Couchbase's SQL++ query engine or Cappella iQ (using natural language) to query the data. Analytical functions can also be written back to the operational side of Capella, where they can be embedded in operational-based dashboards or applications.

The company also released the Capella Columnar SDK for developers. Currently supporting Python, Java and JavaScript, the new SDK reduces the friction of working with columnar data within Capella. The SDK handles all the data connections, including maintaining a database topology and incorporating failover capabilities. Further, SQL++ is supported and can handle both buffered and streaming reads.

Targeted primarily at a developer audience, Couchbase released a suite of four AI services (in preview) for handling unstructured data, vectorizing data, using AI models (both embedding models and large language models [LLMs]) and an agent catalog. The Unstructured Data Service is a pre-processing step that prepares unstructured data such as PDFs for vectorization, and it provides users with options on how to chunk up their unstructured data — for example, by sentence, paragraph or full document, and whether or not to include images. For the Vectorization Service, users can choose which embedding model to use to convert chunked data to vectors. For the Model Service, users can choose from Mistral or Llama 3, which are currently supported. Additional supported models are on the road map.

Couchbase envisions its customers using AI services for RAG (retrieval-augmented generation) scenarios. 451 Research's Voice of the Enterprise: Data & Analytics, Data Architecture for AI 2024 survey reveals that customers source their RAG pipelines from multiple data sources to drive better quality generated responses. With Couchbase's foray into expanding Capella to

store multiple data types, sourcing data from multiple external systems could be minimized.

Sourcing data for RAG

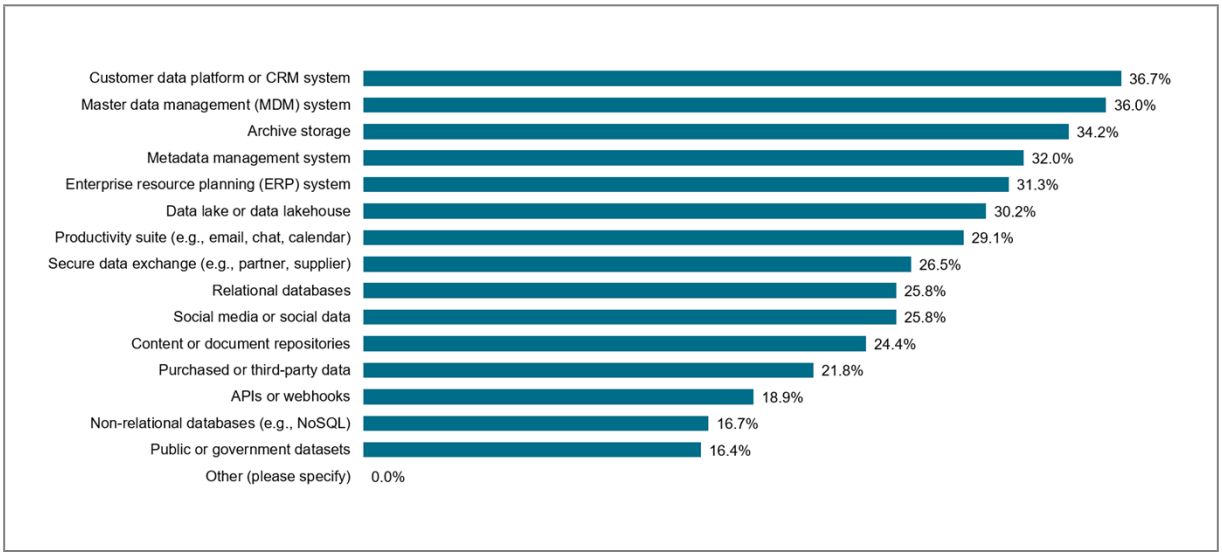

Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Architecture for AI, 2024.
Q. In performing retrieval-augmented generation (RAG), where does your organization primarily source (or plan to source) the "external data" necessary to augment the original LLM training data? Please select all that apply.
Base: Organizations using retrieval-augmented generation (RAG) techniques or plan to in the next 12 months (n=275).

Last, the Agent Catalog Service acts as a type of centralized repository of all the components needed for building and deploying AI agents. The underlying benefit is that it gives users access to the underlying agentic parts to both build or maintain agents.

Competition

Couchbase's primary competition comes from other NoSQL database providers, including MongoDB Inc., DataStax, Redis Labs and MarkLogic (part of Progress), as well as NoSQL database cloud services from the likes of Amazon Web Services, Google, Microsoft Corp. (Azure), Alibaba Cloud, IBM Corp. and Oracle Corp.

Operational systems and analytics systems traditionally have been separate systems, although today many operational database systems will do both within a single environment — what we refer to as hybrid operational and analytical processing. The difference between many hybrid systems lies in the level of analytical capabilities these operational systems possess. A few vendors worth noting include Oracle, IBM, Microsoft, Google, SAP SE, SingleStore, Actian, InterSystems, FairCom, PingCAP, MariaDB and Percona. There are also several NoSQL vendors, such as DataStax, MongoDB, Aerospike and Redis, all of which can accommodate hybrid workloads to some degree.

The introduction of Capella AI Services may put Couchbase in more direct competition with so-called RAG-as-a-service providers, especially because many of these vendors require database integration to some degree. Vendors in this space include Dataworkz, Vectara, Coveo, Nuclia, Contextual, Cody, Valprovia, Ragu, PostgresML, Dezzai and Ragie.

SWOT Analysis

Strengths	Weaknesses
Capella is now a fully multi-modal and multi-model platform, handling transactional, analytical, mobile and AI workloads, all within a single environment with a unified interface. The company's mature portfolio of products includes a managed cloud, on-premises and mobile edge database, as well as an SQL-like query tool for those transitioning to a nonrelational environment.	The newly added analytics service targeted at JSON data is still in its early stages and not yet ready to compete with traditional data warehousing systems, but it is based on columnar storage architecture and has the foundational bones to eventually handle enterprise analytics.
Opportunities	**Threats**
The platforming of Capella gives Couchbase customers and future customers a one-stop computing environment that negates the need to add separate stand-alone tools for RAG scenarios and analytics that cater to JSON data.	The NoSQL database sector remains crowded, with multiple vendors vying for what is nearly one-quarter of the overall database market dominated by industry heavyweights. This is expected to increase even more given the heightened interest in generative AI.

Source: 451 Research.