The consumer of next gen data infrastructures is... *AI*

Capgemini | MongoDB

AI is transforming business, which means a new generation of applications and data infrastructure, designed for AI. The next generation of business applications will serve AI as a business user, creating many new opportunities. But enterprises must re-architect their applications to be AI-native – and that means a new generation of data infrastructures which see AI and AI applications as its primary consumer.

Businesses across all sectors already recognize the value of AI – and ever since ChatGPT burst on the scene in late 2022, a majority of global enterprises have made leveraging generative AI a priority. Today, these powerful technologies are revolutionizing business applications as AI enables better, more personalized customer service, makes recommendations to online shoppers, assists business analysts with anomaly detection, and more.

But this also presents organizations with a huge challenge, as their existing business applications were not built to accommodate AI as a business user. To realize the full competitive advantages of AI, companies need to rebuild their business applications to be AI-native – optimizing them for data accuracy and bringing together structured and unstructured data in a unified database. It's a significant undertaking – but deploying AI itself can ease the transition.

The benefits of data mastery

The good news is that many companies are making great strides towards becoming what Capgemini calls Data Masters – and are already reaping the benefits. In its 2024 report, *Data-powered enterprises: The path to data mastery,* the Capgemini Research Institute noted significant progress since publishing its previous report in 2020. For example, 65 percent of executives surveyed said their organization is harnessing activated data to introduce new products or services – up from 42 percent in 2020. Meantime, 66 percent reported they're leveraging data to launch new business models – compared to 39 percent four years previously.

The Capgemini Research Institute also noted the business opportunity benefits continue to grow for data-powered enterprises. In preparing its report, the institute determined the companies that have mastered their data are significantly outperforming their peers – delivering an average of 39 percent better margins, 42 percent better growth, and 28 percent better earnings before interest and taxes.

AI is the new priority, with new challenges

Capgemini's report observed many of the data mastery challenges faced four years ago have become less worrisome as enterprises pivot to addressing new priorities such as harnessing the power of generative AI. Here too, progress is being made. For example, 60 percent of organizations said they have implemented generative AI pilots or launched early, proof of concept generative AI initiatives based on enterprise data.

However, 75 percent of data executives surveyed admitted large-scale deployment of generative AI continues to be a major challenge. Only 42 percent said they have the required data to train generative AI models, while just 40 percent reported their organizations have developed and deployed the required non-technical foundations – such as corporate culture, ethical guardrails, governance mechanisms, and legal and regulatory frameworks.

The need to address this is becoming more urgent as **the value of data shifts – from being used to report on what has happened, to being used to drive what can happen**. AI is powering this shift, and the technology is becoming easier to use – thanks to changes in the developer stack, where two technologies have enabled applications developers to create AI solutions in ways not previously possible:

- The first is Large Language Models, which power the conversational, natural language interfaces that enable developers to create solutions such as chat bots and agents.
- The second is vector databases, which enable systems to encode unstructured data and then compare them for similarity. This unstructured data – including images, audio, and video – accounts for approximately 80 percent of information in the world today.

These two technologies have helped move AI beyond the exclusive realm of data scientists. Today, applications developers with very little specialized knowledge can build a viable, AI-driven solution by drawing upon open-source hosting and models, and using APIs to access powerful AI infrastructure. These solutions tend to fall into one of four broad categories – semantic search, retrieval augmented generation, recommendations, and anomaly detection – and generally work on live data that's in the end-users' hands even as they're interacting with the application.

Data in silos hinders innovation

These new applications are only going to be as good as the data available to them – and for many enterprises that's now becoming a significant hurdle. Most companies do not have the required data systems in place to support what multi-cloud developer data platform provider *MongoDB* calls operational, as opposed to back office, AI. Data management was cited by 75 percent of developers surveyed as the most challenging part of building AI-powered applications. Meantime, 87 percent of CIOs admitted the ability to bring data to AI and then act upon it is the biggest obstacle they must address.

The cause of this challenge is clear: 59 percent of C-suite executives confessed their enterprise data is either somewhat or completely fragmented into silos – making it inaccessible across the organization.

The lack of cohesion has limited the ability of enterprises to innovate, resulting in lost opportunities: instead of focusing on building cutting-edge applications, teams devote significant time and resources to synchronize disparate systems and manage data complexity. MongoDB simplifies your data architecture and cuts costs by eliminating the need for bolt-on technologies. With native support for many use cases – including semantic and similarity search, geospatial search, stream processing, and time series collections – MongoDB delivers more value at a lower cost per operation.

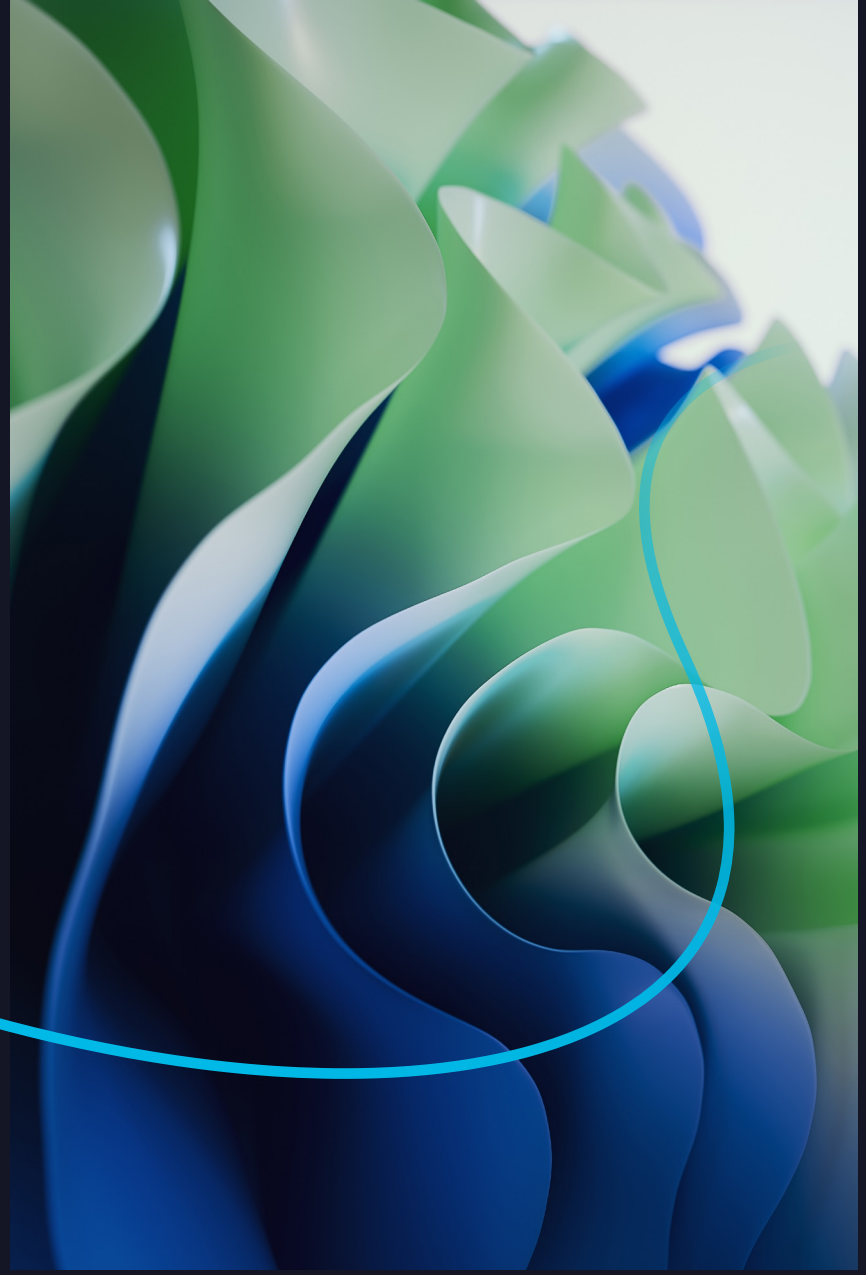

From back office to business user

To address this, enterprises must recognize that, as the role of AI changes, the infrastructure that supports it must also change. Traditional AI infrastructures were designed for backend analytics, where data scientists used complex tools to extract insights from data warehouses. This was effective for reporting and research but prevented AI from being applied to the operational applications that business users interact with daily.

Deploying AI into these end-user applications requires new approaches. For example, while data quality remains important, business users can't wait for data scientists to clean up and verify information after it's collected. Applications must now do this proactively – requiring companies to redesign their applications to ensure they can control data accuracy.

Picking the platform

Marrying operational data with AI infrastructure in a unified database is an effective way to do this. In the opinion of Capgemini and MongoDB, this is no longer optional but fundamental to building AI-powered applications that can scale and adapt to modern business demands. The right data platform is essential – and decision-makers should embrace four pieces of advice when choosing this platform.

1) Don't separate the operational database from the vector database. Keeping them in the same data store means enterprises don't have to synchronize data across multiple systems. It also means developers don't have to query multiple systems and then integrate the results within the application.

2) Choose a data platform that supports a unified query language. This enables developers to easily draw upon diverse workloads – for example, geospatial data, structured data, and vector data – to create compelling applications.

3) Choose a system that's distributed and serverless to ease the demands that diverse workload types place upon hardware.

4) Align your AI application to your business operation model, you don't want unclear business ownership of AI.

The ideal solution is a single logical database, but one that's physically distributed and deploys hardware that's specialized for each type of workload.

MongoDB's document model leads to simpler data models, faster development, and better price per performance for various applications – including systems of engagement like e-commerce platforms, systems of record like inventory management, and systems of insight like Customer 360. MongoDB is also the ideal database for AI workloads because it can handle any kind of data, including complex data types like vector embeddings for images, video, audio, text, and more – all of which are essential for AI models. Because MongoDB's document model and distributed architecture are designed from the ground up to handle any type of data at scale as new technologies become available, MongoDB is an ideal foundation for the AI-powered applications of today and tomorrow.

AI assists in re-architecting applications

Redesigning applications to optimize them for AI is a huge undertaking. Organizations typically have significant legacy systems and a large installed base of applications that worked well in the past but were not designed for the AI world.

Fortunately, artificial intelligence can help with this. As an example, Capgemini has used generative AI to learn how a client's systems work and then define data-migration scripts. This enabled Capgemini to re-architect the client's applications for a new destination on MongoDB – transforming the systems from legacy databases that considered only transactions to a new architecture that supports deploying and managing AI.

The solution leverages Capgemini's Augmented Software Product Engineering Accelerator – an agentic framework to learn, test, verify, and migrate applications. Rather than translating code, this accelerator actually shifts a program from one architecture to another.

Use case demonstrates benefits

With the right data platform, enterprises can leverage AI to create new applications that improve operations. In a recent, real-world example, MongoDB and Capgemini helped a global energy company transform an essential application to take advantage of generative AI.

The company employs field workers in hazardous conditions. In the past, these workers were required to fill out forms manually, while on the jobsite. In some cases, this involved forms with approximately 75 fields. This was a complicated process that caused issues when working, for example, on an oil rig – where employee safety is paramount.

Capgemini and MongoDB re-architected the application to incorporate generative AI and natural language processing. The new solution uses a backend that combines an operational database and a vector store, and has been deployed via Azure Cloud. Today, instead of filling out forms manually, workers provide the necessary information via a conversational interface. This has been so successful that the client has been deploying the application to 120,000 workers around the globe, which has dramatically increased safety and productivity.

Focus on data infrastructure

The evolution of AI into the operational activities is an exciting transformation – but it requires enterprises to start by building better operational data infrastructures. It's the only way organizations can ensure they're in control of their data and can deliver the right information to the AI, at the right time. Companies that do it right will be better able to capitalize on the power of these technologies that create significant business opportunities and competitive advantages.

Authors

Please reach out with questions or to schedule a conversation about this paper's content and our capabilities to assist your organization:

Steve Jones

Executive Vice President - Collaborative Data Ecosystems, Capgemini

Prasad Pillalamarri

Director for SI Partner Solution Consulting, MongoDB

About Mongo DB

Headquartered in New York, MongoDB's mission is to empower innovators to create, transform, and disrupt industries by unleashing the power of software and data. Built by developers, for developers, our developer data platform is a database with an integrated set of related services that allow development teams to address the growing requirements for today's wide variety of modern applications, all in a unified and consistent user experience. MongoDB has tens of thousands of customers in over 100 countries. The MongoDB database platform has been downloaded hundreds of millions of times since 2007, and there have been millions of builders trained through MongoDB University courses. To learn more, visit *mongodb.com.*

MongoDB: The Developer Data Platform

Get your ideas to market faster with a developer data platform built on the leading modern database. MongoDB makes working with data easy.

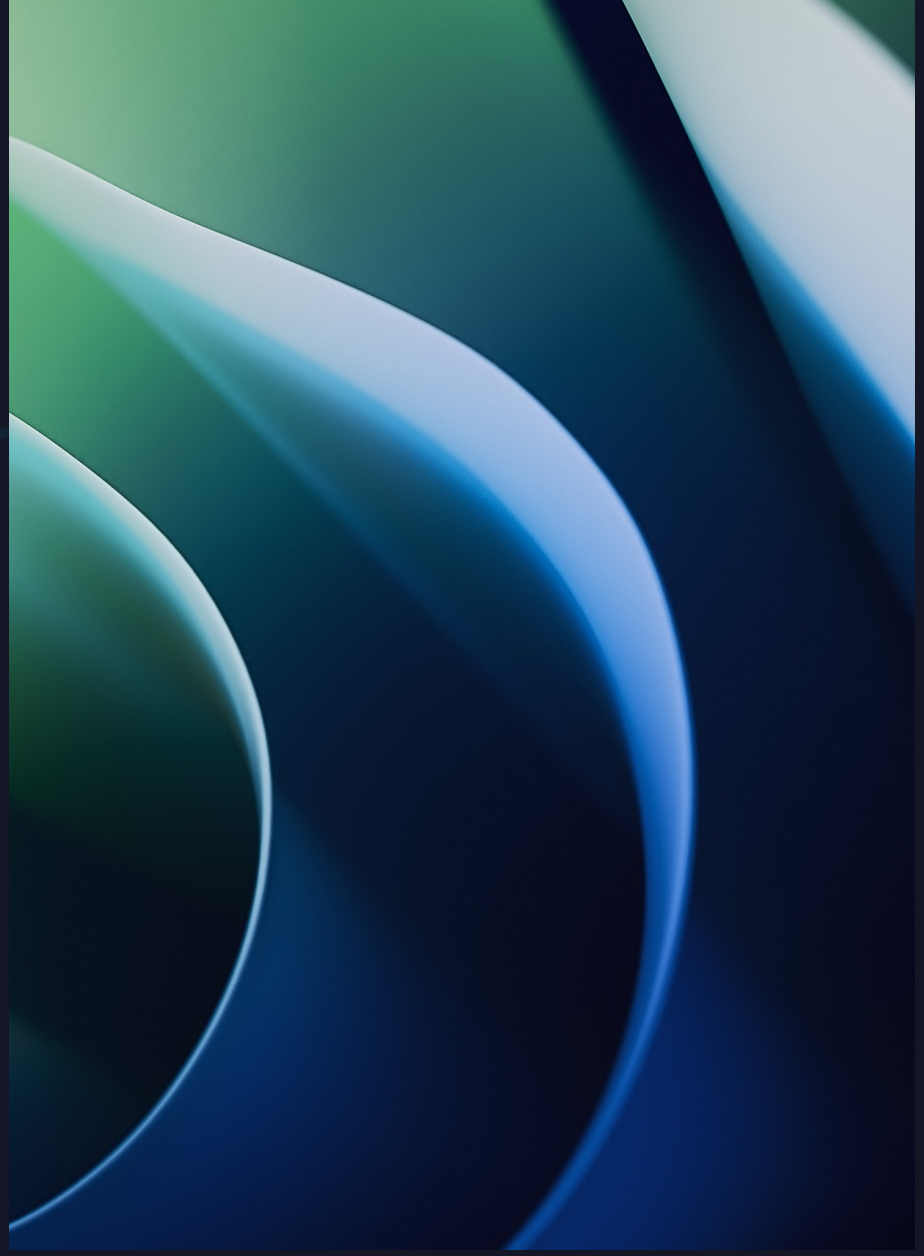

About Capgemini

Capgemini is a global business and technology transformation partner, helping organizations to accelerate their dual transition to a digital and sustainable world, while creating tangible impact for enterprises and society. It is a responsible and diverse group of 340,000 team members in more than 50 countries. With its strong over 55-year heritage, Capgemini is trusted by its clients to unlock the value of technology to address the entire breadth of their business needs. It delivers end-to-end services and solutions leveraging strengths from strategy and design to engineering, all fueled by its market leading capabilities in AI, cloud and data, combined with its deep industry expertise and partner ecosystem. The Group reported 2023 global revenues of €22.5 billion.

Get the future you want | www.capgemini.com