

Enterprises use multicloud to power agentic AI and next-generation applications, supporting global scalability, regulatory compliance, and cost efficiency while accelerating digital modernization.

Enabling Next-Generation Applications with Multicloud Platforms

September 2025

Written by: Devin Pratt, Research Director, Data Management, AI, Automation, and Data Practice

Introduction

The rapid evolution of AI is reshaping enterprise IT strategies, with organizations increasingly adopting AI-driven technologies to enhance operational efficiency, decision-making, and innovation. In IDC's June 2025 *Future Enterprise Resiliency and Spending Survey, Wave 5*, 92% of the 885 organizations surveyed reported having a strategy for adopting and utilizing AI technologies in their operations. Nearly all organizations are investing in or embedding generative AI (GenAI) across their operations, and only 1.7% are not yet doing the same with AI agents, according to the 893 organizations surveyed in IDC's April 2025 *Future Enterprise Resiliency and Spending Survey, Wave 3*.

As AI workloads expand, the demand for strong, scalable, and low-latency infrastructures has grown significantly. Multicloud environments have emerged as a critical enabler for modern applications, offering organizations the ability to utilize multiple cloud providers concurrently to optimize data strategies, scalability, and cost efficiency. *Worldwide Database Management Systems Software Market Shares, 2024: Powering AI Workloads and Enabling Autonomous DBMS Operations* (IDC #US53032425, June 2025) highlights that public cloud services accounted for 62.2% of database management system revenue in 2024. This highlights the increasing migration of workloads to cloud environments.

However, the adoption of multicloud architectures presents challenges, particularly when seamless data connectivity across regions and clouds is essential. AI workloads often require real-time analytics and low latency, which demand sophisticated orchestration and integration capabilities. Unified data strategies that integrate structured and unstructured data across multiple clouds are essential for enabling comprehensive data utilization and real-time, accurate decision-making.

This paper delves into the intersection of multicloud platforms and agentic AI, exploring key trends, benefits, and challenges. It also examines the role of unified data strategies in driving modernization and innovation.

AT A GLANCE

KEY STAT

According to IDC, by 2028, there will be 1.3 billion agentic AI agents worldwide; 92% of enterprises already have an AI strategy in place.

KEY TAKEAWAY

To meet rapidly growing AI demands, enterprises need multicloud architectures that support scalability, flexibility, low latency, and cost efficiency.

Unified Architecture for Agentic AI and Multicloud Transformation

By 2028, IDC predicts 1.3 billion AI agents will be in operation globally, transforming workflows and decision-making processes. The next phase will prioritize orchestrating interactions among agents, ensuring reliable tool integration, and establishing safeguards for autonomous operation across business functions. This shift elevates requirements for governed access to operational data, consistent policy controls, and observability across clouds, capabilities that multicloud and unified data strategies are designed to deliver. Elements of this shift include the following key capabilities:

- » **AI-ready infrastructure.** Enterprises are investing in infrastructure designed to meet the demands of agentic AI. An IDC United States–based survey of 102 organizations found that 74.5% of organizations plan to adopt integrated vector databases to store and query embeddings in their agentic AI workflows (see Figure 1). This illuminates a broader shift toward embedding vector search directly into existing data platforms. The aim is reduced data movement, streamlined governance, and lower latency for AI applications.
- » **Unified data strategies (including hybrid transactional/analytical processing).** Integrating structured and unstructured data across multiple clouds enhances AI performance by enabling comprehensive data utilization and more accurate decision-making. The convergence of online analytical processing (OLAP) and online transaction processing (OLTP) data, achieved through hybrid transactional/analytical processing (HTAP), supports agentic AI in two key ways. First, it reduces batch delays. Second, it enables real-time context for planning, recommendations, and autonomous actions.

In IDC's June 2025 *Data Culture and Decision Intelligence Survey*, 764 organizations highlighted two primary areas for advancing their decision-making intelligence solutions: improving predictive accuracy (51.1%) and enhancing real-time processing capabilities (47.0%). In practice, these priorities translate into shared metadata and governance, streaming pipelines, and consistent policy enforcement across clouds.


- » **Multicloud integration.** Organizations are adopting multicloud architectures to deliver efficiency, control, and regulatory alignment at scale. Of the 675 organizations that responded to IDC's February 2025 *Enterprise Infrastructure Survey*, 60.1% said cloud service providers were the preferred channel for executing either a hybrid or multicloud solution. Done right, multicloud becomes a foundation for agility, resilience, and long-term competitiveness.
- » **Data sovereignty and compliance.** By ensuring compliance with regulations such as GDPR, data sovereignty enables organizations to store data in specified regions. According to IDC's 2025 *Digital Sovereignty Survey*, 36.8% of the 955 organizations surveyed use sovereign public cloud solutions, with an additional 43.0% planning to adopt them. Among those using sovereign cloud infrastructure, 44.5% rely on public cloud platforms provided by global cloud service providers augmented with sovereign controls for security and compliance. Sovereign cloud adoption ensures organizations meet regulatory requirements while maintaining trust and control over sensitive data.
- » **Interoperability and flexibility.** Multicloud environments help organizations avoid vendor lock-in while providing a combination of scalability, performance, and cost advantages:
 - **Improved scalability.** Multicloud architectures enable global application scaling by employing the strengths of multiple cloud service providers. Public cloud services grew by 12.8% in 2024, emphasizing their scalability benefits (see *Worldwide Database Management Systems Software Market Shares*,

2024: Powering AI Workloads and Enabling Autonomous DBMS Operations, IDC #US53032425, June 2025).

- **Low latency.** Reducing the distance between data, compute, and users is essential for agentic AI, where every millisecond matters for real-time retrieval and inference.
- **Cost efficiency.** Portable workloads enable enterprises to balance price performance across providers, minimize overhead from egress and replication, and dynamically size infrastructure to evolving utilization patterns.

By uniting these capabilities, enterprises gain an architecture that powers responsible AI at scale, unlocks data value, and accelerates business impact with confidence.

FIGURE 1: *Infrastructure for Agentic AI Workflows*



n = 102

Source: IDC's U.S. IT Quick Poll — Agentic AI and Data Survey, 2025

Considerations

As enterprises scale agentic AI across multicloud environments, they face complex challenges that span connectivity, compliance, and cost control:

- » **Consistent, low-latency data connectivity and integration.** Delivering reliable, low-latency data connectivity across multiple cloud providers requires advanced orchestration and networking to support real-time AI workloads. Enterprises must coordinate data flows across clouds, manage event streaming, and align API policies. Concurrently, they must address challenges such as data gravity, egress costs, and metadata consistency. To ensure smooth operations, organizations need end-to-end service-level objectives (SLOs) and observability across regions and providers. Without these measures, AI agents risk working with stale context, producing inconsistent results, or suffering performance slowdowns during traffic spikes.

- » **Data sovereignty and compliance at scale.** Meeting diverse sovereignty and compliance requirements require precise residency controls and governance frameworks tailored to each region, ensuring adherence to GDPR and other local regulations. As AI agents gather context and generate outputs, organizations need location-aware data routing, regional key management, auditable data lineage, and policy as code that enforces locality from ingestion through inference and logging. Cross-border transfer restrictions must be addressed, and "right to be forgotten" processes implemented. Organizations must also comply with sector-specific rules in areas like financial services and the public sector. All of this must be done without breaking the cohesion of the overall data strategy.
- » **Balancing dynamic scaling with cost efficiency.** Managing dynamic workload scaling while controlling costs requires careful resource allocation and workload placement across clouds. This minimizes duplication, avoids vendor lock-in, and keeps spending under control. AI pipelines — covering data preparation, vectorization, retrieval augmented generation (RAG), fine-tuning of foundation models, domain-specific adaptation, prompt optimization, and inference — create fluctuating demand for compute, storage, networking, and GPUs across regions.

Success depends on embedding compliance, observability, and cost-aware scaling into multicloud strategies to support the reliability and responsiveness agentic AI requires.

Considering MongoDB Atlas

MongoDB Atlas is positioned as a multicloud data platform intended to support agentic AI and real-time, data-intensive applications. It is designed to unify structured and unstructured data in a JSON document model with native vector search and in-application, real-time analytics on operational data (HTAP style). It also supports native cross-cloud deployment and cluster distribution across AWS, Microsoft Azure, and Google Cloud.

The platform's features and capabilities include:

- » **AI-ready infrastructure.** Atlas is presented as an integrated vector database approach, designed to enable native vector search alongside JSON document storage. This allows organizations to handle diverse AI workload requirements and implement agentic AI workflows with fewer data hops.
- » **Low-latency operations.** The platform is optimized for low-latency environments and real-time data availability for AI applications that require timely retrieval, inference, and event processing.
- » **Unified data strategy for real-time analytics.** Atlas is intended to unify structured and unstructured data in a JSON document model to power AI and in-application, real-time OLAP on live operational data. Engineering teams can deliver low-latency dashboards, personalization, and anomaly detection without heavy ETL pipelines or long-running batch queries by bringing OLTP and real-time OLAP together in one platform. This approach complements existing traditional warehouses that continue to serve historical and trend analysis use cases.
- » **Multicloud integration.** Atlas supports native deployment across AWS, Microsoft Azure, and Google Cloud, with options for a single cluster's distribution to span multiple providers simultaneously. According to IDC's 2025 *Digital Sovereignty Survey* of 955 organizations, this approach aligns with enterprise preferences for sovereign cloud infrastructure and platform providers.

- » **Global availability and portability.** Atlas is available in 125 cloud regions globally (with ongoing expansion). Optionality for seamless data replication and migration across regions and providers to support modernization initiatives is an inherent design element.
- » **Data sovereignty and compliance.** Zone-based partitioning, multiregion deployment options, and residency controls are designed to help organizations meet data residency requirements and assist with GDPR and regional regulatory needs.
- » **Low risk of cloud provider lock-in.** By using a multicloud database architecture that supports both native vector search and JSON document storage, teams can reduce dependence on any single provider. This approach allows data and workloads to be placed flexibly across clouds while maintaining a unified operational strategy for governance, security, and performance.

As a modern multicloud platform, MongoDB Atlas enables real-time AI, scales globally, and meets regulatory demands in a single solution.

Challenges

As enterprises accelerate AI-driven transformation, MongoDB faces the challenge of ensuring Atlas is clearly differentiated in an increasingly crowded multicloud and data platform landscape:

- » **Ecosystem competition and differentiation.** While Atlas provides multicloud flexibility across AWS, Microsoft Azure, and Google Cloud, hyperscalers themselves continue to expand native database and AI capabilities. These competing offerings can blur differentiation and complicate decision-making for enterprises already embedded in a single cloud provider's ecosystem. MongoDB's ability to emphasize Atlas' native cross-cloud portability, sovereignty controls, and vector-native design will be critical to standing out.

Ultimately, MongoDB's ability to demonstrate that Atlas delivers differentiated value beyond native hyperscaler services will define its success in sustaining long-term relevance in the AI and multicloud era.

Conclusion

The convergence of multicloud architectures and agentic AI is redefining enterprise IT, enabling scalable, real-time, and automated operations. This unified fabric allows agents to act on distributed data with intelligence. Organizations can preserve the agility to deploy workloads where performance and cost align best.

A unified data strategy, integrating both structured and unstructured data across multiple clouds, underpins successful AI performance. It supports HTAP, enabling faster and more accurate decision-making. Establishing shared metadata, governance, and security controls across providers further enhances the reliability of agentic workflows by ensuring consistent policies and access. When combined with in-place data access strategies, these measures reduce unnecessary data movement and can help lower latency.

At the same time, data sovereignty and compliance must remain top priorities. Region-specific controls and sovereign cloud solutions are essential for adhering to GDPR and other international regulations in global deployments.

Reliable low-latency connectivity across multiple providers requires advanced orchestration and network design to effectively support real-time AI workloads. Similarly, dynamic workload placement, balanced with cost optimization, helps organizations avoid vendor lock-in, maximize utilization, and control spending in complex multicloud environments.

For enterprises, multicloud is the key to the success of agentic AI and next-generation applications, delivering portability, scale, and compliance.

About the Analyst



Devin Pratt, Research Director, Data Management, AI, Automation, and Data Practice

Devin Pratt is research director for Data Management in the AI, Automation, and Data practice at IDC, with expertise in data management software and the evolving data plane that supports agentic AI.

MESSAGE FROM THE SPONSOR

Headquartered in New York, MongoDB's mission is to empower innovators to create, transform, and disrupt industries with software. MongoDB's unified database platform powers the next generation of applications and is the most widely available, globally distributed database on the market. With integrated capabilities for operational data, search, real-time analytics, and AI-powered data retrieval, MongoDB helps organizations move faster, innovate more efficiently, and simplify complex architectures. Millions of developers and more than 50,000 customers — including 75% of the Fortune 100, rely on MongoDB for their most critical applications.

The MongoDB Application Modernization Platform (AMP) combines AI-powered tooling, proven methodologies, and engineering expertise to help enterprises transform legacy applications into modern, scalable services. By addressing technical debt and accelerating modernization faster than traditional approaches, AMP enables organizations to redirect resources toward innovation and to capitalize on the promise of AI. With the flexibility of MongoDB as its foundation, AMP helps organizations speed time-to-market, reduce technical debt, and future-proof their businesses.

Learn more at <https://www.mongodb.com/solutions/use-cases/modernize>.



The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. This IDC material is licensed for external use, and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives.

©2025 IDC. Reproduction is forbidden unless authorized. All rights reserved. [CCPA](#)