IDC MarketScape

# IDC MarketScape: Worldwide Operational Databases 2026 Vendor Assessment

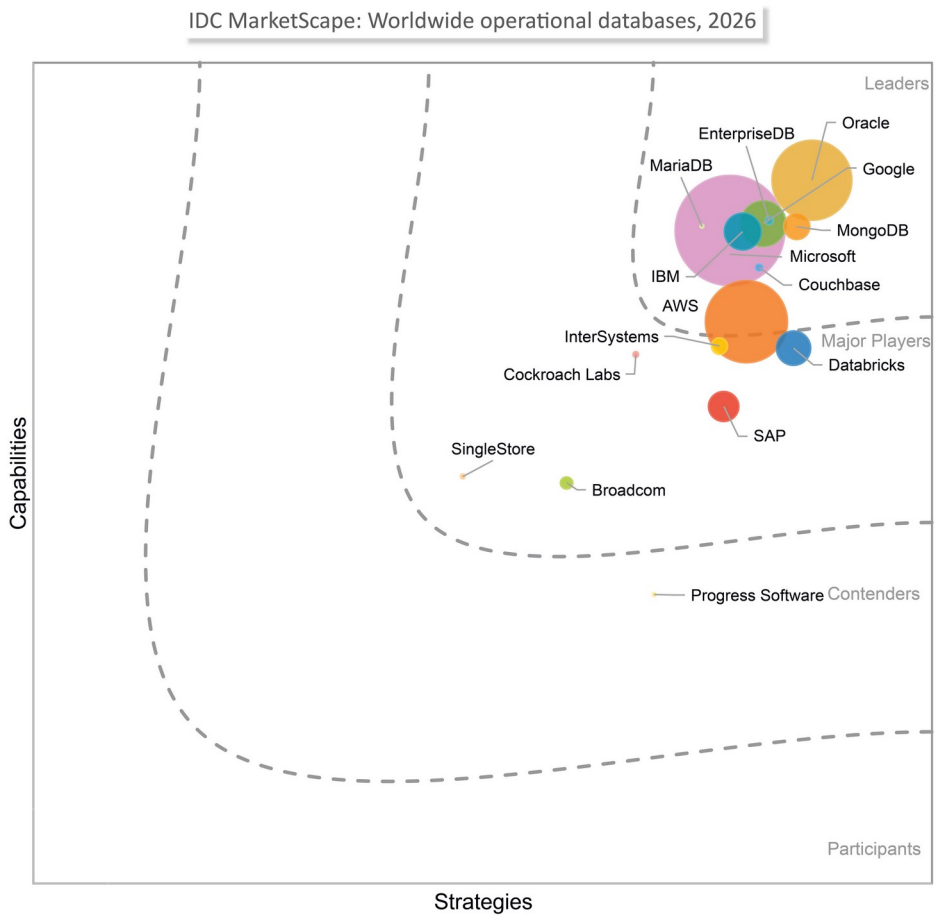
Devin Pratt

## THIS EXCERPT FEATURES MARIADB AS A LEADER

## IDC MARKETSCAPE FIGURE

### FIGURE 1

**IDC MarketScape: Worldwide operational databases vendor assessment**




Source: IDC, 2026

November 2026, IDC #US54124326e

Please see the Appendix for detailed methodology, market definition, and scoring criteria.

## ABOUT THIS EXCERPT

The content for this excerpt was taken directly from IDC MarketScape: Worldwide Operational Databases 2026 Vendor Assessment (Doc # US54124326).

## IDC OPINION

The operational database market is entering a period of rapid change as enterprises balance the demands of always-on transactional systems with the growing need to ground AI and agentic applications in live business data. For technology buyers, this research provides a strategic framework to evaluate operational databases on the capabilities that now separate leaders from the field. Those capabilities are transactional performance and resilience; deployment and interoperability flexibility; and readiness for the vector, semantic, and agentic workloads reshaping enterprise data strategy. The broader database management systems market grew to $107.8 billion in 2025, and IDC forecasts it to reach $192.8 billion by 2030 (see *Worldwide Database Management Systems Software Forecast, 2026–2030,* IDC #US53749326, September 2026). For buyers, though, the more consequential shift is qualitative, not quantitative: Operational databases are being asked to support agentic AI, vector search, and real-time analytics workloads that didn't factor into most purchasing decisions two years ago.

Evaluating those capabilities fairly means not defaulting to company size as a signal of quality. Company scale is a recurring concern buyers raise about independent vendors, but it does not consistently track with measured capability. Offerings from several smaller vendors in this assessment perform at an above-average level for this market. For buyers, that gap between perceived risk and demonstrated platform capability is worth weighing on its own, separate from company size.

Enterprises are moving quickly to converge operational and analytical processing on shared infrastructure. According to IDC's *Worldwide Data Management Survey,* 2026 (n = 715), 82% of organizations rate supporting converged operational and analytical workloads as very or extremely important over the next 12–24 months, and 71% report they have already implemented converged workloads in production or standardized them as an enterprise-wide approach. Yet convergence has not settled on a single architectural pattern. Rather than standardizing on one vendor or one engine, the dominant model emerging across the market is a core platform complemented by selected point solutions connected through a federation layer under shared governance. Only 18% of organizations surveyed expect to run operational and analytical workloads on one primary database going forward, while a combined 53% expect to rely on multiple specialized databases connected through

federation or a unified platform with shared governance. This gap between rising demand for convergence and the reality of continued architectural plurality underscores why architectural flexibility, rather than a single winning pattern, is the defining characteristic of this market today.

Agentic AI is the strongest force behind this change, and it is rewriting who buys the database. Thirty percent of organizations already run agentic AI in production, and another 50% are investing with an established spending plan (source: IDC's *Worldwide Data Management Survey,* 2026; n = 715). As autonomous agents become active consumers of enterprise data rather than just applications, vendors are increasingly designing for three distinct personas: the executive buyer focused on consolidation and ROI, the DBA managing operational efficiency and autonomous operations, and the agent itself as an autonomous consumer of data requiring governed, auditable access. Operational databases are becoming a primary home for the vector data these agents depend on. Sixty-three percent of organizations use or plan to use an integrated vector database, such as PostgreSQL with the pgvector extension, for agentic AI workflows, on par with the 63% relying on data warehouse or lakehouse platforms for the same purpose. Vector embeddings alone are not enough for agents to reason reliably about enterprise data. IDC's February 2026 *U.S. Data Management Quick Poll Survey* found that 47% of respondents identify a consistent semantic layer, spanning ontologies, knowledge graphs, and agent memory, as critical to safe and trustworthy AI. That data readiness does not extend to how databases run themselves: Applying AI to tuning, anomaly detection, and self-healing remains an area where vendors more often differ. As agentic AI expands the data these teams manage and staffing stays lean, this kind of automation frees capacity for higher-value work rather than routine maintenance.

Deployment reality remains stubbornly hybrid, fragmented, and sprawling. Just 14% of organizations describe their operational database environment as fully standardized on one primary technology, and organizations now manage an average of 16.9 distinct database technologies alongside 16.8 additional data platforms and management tools, before any consolidation work begins (see *IDC Survey: Enterprise Buyer Behavior and Data Management Consolidation Trends,* IDC #US54555726, July 2026; source: IDC's *Worldwide Data Management Survey,* 2026; n = 715). That sprawl is precisely why 97% of organizations are actively pursuing some degree of consolidation across data management capabilities over the next 24 months, even as open source adoption becomes nearly universal, with 90% of organizations running at least some open source operational databases in production. PostgreSQL in particular has become the foundation of choice for this shift, with hyperscalers and independent vendors alike building cloud, Kubernetes, and AI-native offerings on top of it, even as established commercial engines retain the deepest enterprise installed bases. Data sovereignty compounds the deployment challenge. Approximately 95% of organizations report that sovereignty and residency requirements influence where they can store or process data, pushing enterprises toward hybrid and

multicloud architectures rather than a single public cloud. Taken together, sprawl, consolidation pressure, open source adoption, and sovereignty mean the operational database must earn its place not as a standalone system but as one governed, portable component of an estate that spans engines, clouds, and jurisdictions.

The operational database is no longer a passive system of record. It is becoming the control point where retrieval, governance, and current context are enforced, and the market's fastest-moving vendors are engineering for that role directly. Leadership in this market will belong to platforms that pair transactional trust with the openness, deployment flexibility, and embedded AI that enterprises are already converging toward. The vendors evaluated in this IDC MarketScape demonstrate these capabilities in different combinations, and the profiles that follow are intended to help buyers match those strengths to their own requirements.

## IDC MARKETSCAPE VENDOR INCLUSION CRITERIA

A critical point in this research effort is to meet the following inclusion criteria:

- **Core offering:** The vendor must offer an operational database solution as a core product. This solution should be purpose-built to support transactional workloads and live application data rather than analytical-only workloads.
- **Transactional performance:** The database must support low-latency read and write operations with strong consistency and transaction integrity to power day-to-day business operations.
- **Scalability:** The solution must demonstrate the ability to scale horizontally (e.g., by adding nodes) or vertically (e.g., by increasing resources on a single node) without performance degradation.
- **High availability and resilience:** The solution must provide high availability with automated failover, replication, backup, and disaster recovery capabilities to support mission-critical applications.
- **Concurrency and throughput:** The database must handle high levels of concurrent transactions and users effectively.
- **Deployment flexibility:** The database must offer flexible deployment options across on-premises, public cloud, hybrid, and/or multicloud environments.
- **Global presence:** Vendors must generate at least 10% of recognized revenue from their operational database product in each of the Americas and EMEA and at least 5% in APJ.
- **Revenue threshold:** Vendors must have at least $50 million in annual revenue associated with the operational database software for 2025.

- **Vendor IP ownership:** The core components of the operational database (e.g., transaction engine, storage layer, replication, and query processing) must be the intellectual property of the vendor. Partnered or OEM-provided IP is acceptable for supplementary features such as monitoring, data integration, or governance.
- **Strategic importance:** The operational database undergoing evaluation must represent a strategic offering within the vendor's portfolio.

## ADVICE FOR TECHNOLOGY BUYERS

Based on IDC's research and observed market dynamics, organizations evaluating operational database platforms should consider the following guidance:

- **Treat security, privacy, and compliance as the primary selection filter, not an afterthought.** Security, privacy, and compliance are the top 3 most cited purchase criteria in IDC's *Worldwide Data Management Survey,* 2026 (n = 715), named by 49% of organizations selecting or expanding an operational database, ahead of raw performance and scalability (32%) and availability and resilience (28%). Buyers should confirm published SLAs, certification breadth, and recovery time objective/recovery point objective (RTO/RPO) targets alongside performance benchmarks rather than treating security posture as an afterthought to speed and throughput.
- **Plan for a mix of engines, not a single winning database, even as convergence accelerates.** Only 18% of organizations expect to consolidate onto a single primary database, while a combined 53% expect to run multiple specialized engines connected through federation or a shared-governance platform, even as convergence becomes table stakes. The market's emerging pattern is not single-vendor standardization but a core platform complemented by selected point solutions under shared governance. Buyers should weigh a vendor's interoperability, open-format support, and federation capabilities as carefully as its native convergence story and favor platforms that can act as that organizing layer rather than one more point solution to integrate.
- **Treat vector storage, semantic context, and agentic AI enablement as near-term requirements, not future road map items.** Buyers should confirm native or well-integrated vector search, embedding storage and agent-enablement features now rather than planning to retrofit them later. Vector support alone is not sufficient; evaluate a consistent semantic layer that covers ontologies, knowledge graphs, and agent memory as part of the same decision.
- **Evaluate AI-driven operations, not just AI-ready data, when assessing automation.** Native vector and AI-ready data support is now common across the market, but applying AI to the database's own operations (autonomous tuning,

anomaly detection, and self-healing) remains uneven. Buyers should distinguish between vendors offering genuine AI-driven operational intelligence and those still relying on traditional scripted automation since this gap affects how much routine maintenance a leaner team can offload.

- **Size up your organization's real database sprawl before committing to a consolidation path.** Only 14% of organizations describe their environment as fully standardized today, even as consolidation efforts accelerate. Buyers should audit their actual database and tooling footprint first, then weigh a platform's hybrid and multicloud deployment options, migration tooling, and licensing flexibility against the realistic pace of consolidation their organization can sustain.
- **Weigh sovereignty and multicloud portability as architecture requirements, not compliance checkboxes.** Data sovereignty and residency requirements increasingly influence where organizations can store or process operational data, and public cloud database revenue is growing far faster than on-premises revenue marketwide. Buyers operating across borders should confirm native geographic data pinning, local encryption key management, and consistent governance across hybrid, sovereign cloud, and on-premises environments and should be wary of any platform that requires full cloud consolidation to deliver those guarantees.

Source: IDC's *Worldwide Data Management Survey,* 2026 (n = 715) , unless otherwise noted

## VENDOR SUMMARY PROFILES

This section briefly explains IDC's key observations resulting in a vendor's position in the IDC MarketScape. While IDC evaluated every vendor against each of the criteria outlined in the Appendix, the description here provides a summary of each vendor's strengths and challenges.

### MariaDB

After a thorough evaluation of MariaDB's strategies and capabilities, IDC has positioned MariaDB in the Leaders category in this 2026 IDC MarketScape for worldwide operational databases vendor assessment.

MariaDB's operational database offering is MariaDB Enterprise Server (release 11.8), a hardened, fully open source relational database that anchors the MariaDB Enterprise Platform. Available since 2009 as a drop-in successor to MySQL and with deep Oracle compatibility, including native PL/SQL support, it has a broad global installed base across sectors that value an open, standards-based relational database with commercial hardening and support. It handles transactional, analytical, and AI workloads within a single platform and is designed for portability across on-premises, cloud, and hybrid environments.

MariaDB provides a pluggable storage architecture that natively supports row-based, columnar, hybrid, and in-memory engines selectable per table, alongside converged transactional and analytical processing through its ColumnStore engine. It offers synchronous multimaster clustering through Galera for continuous availability, a Kubernetes operator for full life-cycle automation, native vector capabilities for AI, and full transaction-isolation support. Its open source foundation, MySQL wire compatibility, and long-standing Oracle PL/SQL compatibility make it broadly accessible to developers and portable across environments, without requiring applications to be rewritten.

With an open source model that avoids vendor lock-in, core transactional and analytical capabilities, broad developer familiarity, and an open source database community, MariaDB positions itself as a flexible, cost-effective foundation for organizations modernizing relational workloads. This approach aligns well with enterprises seeking an open, portable database that can consolidate diverse workloads while retaining enterprise-grade resilience and commercial support.

## Strengths

### Pluggable storage architecture for diverse workloads

MariaDB natively supports row-based, columnar, hybrid, in-memory, and vector storage architectures, with pluggable engines selectable per table to match each workload to the most suitable storage model. This allows a single database to serve transactional, analytical, mixed, and AI/vector workloads without adopting separate specialized systems, including native vector search for GenAI and retrieval-augmented generation built into the engine rather than added as a separate system. For organizations with varied data-access patterns across their applications, this adaptability supports workload consolidation and lets teams tune storage to performance and cost requirements table by table, rather than compromising on a single fixed engine.

### Converged transactional and analytical workloads

MariaDB supports converged operational and analytical processing within a single platform through its ColumnStore engine and isolates mixed workloads through separate compute, separate replicas, and resource governance. This allows analytical queries to run alongside transactional workloads without one degrading the other and without moving data to a separate analytical system. For organizations seeking to reduce the number of specialized data stores in their architecture, this convergence supports operational reporting and analytics directly against transactional data while preserving performance and consistency. MariaDB's broader Enterprise Platform also offers MariaDB Exa, an additional analytical and HTAP engine for enterprises that need further scalability and performance.

### Full life-cycle automation

MariaDB provides a Kubernetes operator that automates the complete operational life cycle, including provisioning, scaling, upgrades, failover, backups, and restore, supporting consistent management across on-premises and cloud deployments. Combined with synchronous multimaster clustering for continuous availability, this automation reduces the day-to-day administrative effort of operating the database at scale. For organizations weighing the total cost and staffing required to run an open source database in production, this operational tooling brings commercial-grade automation to a platform that retains open source flexibility.

## Challenges

### Smaller company scale

MariaDB is a smaller, privately held company, which can lead some buyers to weigh vendor scale heavily in evaluations. Rather than concentrating on company size alone, organizations should evaluate MariaDB against their own enterprise requirements and the platform's demonstrated capabilities, including its pluggable storage architecture, converged workload support, and resilience since these determine fit more directly than scale.

### Maturity of AI-assisted operations

MariaDB provides self-healing and automatic failover remediation as part of its operational tooling, but its broader AI-assisted operational capabilities, such as autonomous tuning and anomaly detection, are limited. MariaDB's life-cycle automation, self-healing, and alerting already address much of the day-to-day operational burden, and the platform separately enables AI workloads directly through native vector search, integrated retrieval-augmented generation, and MCP server support. Organizations that rely heavily on AI-driven database operations should evaluate the current state of these capabilities against their expectations, though the company continues to expand AI-assisted capabilities as part of its ongoing road map, including AI-assisted migration tooling and an enhanced DBA copilot.

### Breadth of compliance certifications

MariaDB holds a set of compliance attestations, including ISO 27001, SOC 2 Type II, HIPAA, and GDPR, along with stated conformance to frameworks such as STIG, FIPS, and CCPA, but its formal certification coverage is narrower in areas such as PCI DSS and government-specific programs. Because customers frequently deploy MariaDB within their own environment, overall compliance posture depends in part on the surrounding infrastructure. Organizations in regulated industries should evaluate MariaDB against the specific certifications their deployment and regulatory scope require, rather than certification count alone.

### Consider MariaDB when

Consider MariaDB when an open, portable, and cost-effective relational database that avoids vendor lock-in is a priority, particularly for organizations modernizing MySQL-compatible or Oracle-compatible workloads, or when consolidating diverse data-access patterns onto a single platform. The platform suits enterprises that value a pluggable storage architecture, converged transactional and analytical processing, native vector search, and continuous availability through multimaster clustering, all on an open source foundation with commercial hardening and support. Organizations seeking developer familiarity, deployment portability, and workload flexibility will find MariaDB a practical option for balancing openness, capability, and enterprise-grade reliability.

## APPENDIX

## Reading an IDC MarketScape graph

For the purposes of this analysis, IDC divided potential key measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis reflects the vendor's current capabilities and menu of services and how well aligned the vendor is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here, and now. Under this category, IDC analysts will look at how well a vendor is building/delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the x-axis, or strategies axis, indicates how well the vendor's future strategy aligns with what customers will require in three to five years. The strategies category focuses on high-level decisions and underlying assumptions about offerings, customer segments, and business and go-to-market plans for the next three to five years.

The size of the individual vendor markers in the IDC MarketScape represents the market share of each individual vendor within the specific market segment being assessed. For this IDC MarketScape, IDC determined the vendor market size using IDC's Worldwide Semiannual Software Tracker, which represents an estimate of each vendor's 2025 database management systems software revenue.

## IDC MarketScape methodology

IDC MarketScape criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and specific vendors. IDC analysts tailor the range of standard characteristics by which vendors are measured through structured discussions, surveys, and interviews with market leaders, participants, and end users. Market weightings are based on user interviews, buyer surveys, and the input of IDC

experts in each market. IDC analysts base individual vendor scores, and ultimately vendor positions on the IDC MarketScape, on detailed surveys and interviews with the vendors, publicly available information, and end-user experiences in an effort to provide an accurate and consistent assessment of each vendor's characteristics, behavior, and capability.

Within the capabilities axis, native delivery within the database engine had a slightly higher value than functionally equivalent capabilities delivered through a companion product or platform, reflecting this IDC MarketScape's focus on the operational database itself, though both were treated as legitimate ways to deliver a capability. Capability scores reflect each vendor's generally available functionality as of July 1, 2026.

## Market definition

An operational database serves as the transactional backbone of modern applications and digital business processes. Unlike analytical databases, which are optimized for large-scale querying and analysis, operational databases are designed to capture, process, store, and update live data with low latency, high concurrency, and strong consistency to support day-to-day business operations and application workloads. This market includes relational databases as well as nonschematic and multimodel approaches, such as document and key-value databases, enabling organizations to align database architecture with diverse application and data requirements. As enterprise requirements evolve, operational databases are also incorporating greater automation, hybrid and multicloud deployment flexibility, stronger security and governance, and, in some cases, converged transactional and analytical capabilities to support real-time decision-making.

Key features of operational databases include:

- **Transactional performance:** Low-latency, high-throughput read and write operations with strong consistency and transaction integrity to power day-to-day business operations
- **Scalability:** The ability to scale horizontally and vertically, including scaling reads and writes independently, without performance degradation as workloads grow
- **High availability and resilience:** Continuous availability with automated failover, replication, backup, and disaster recovery to keep mission-critical applications running
- **Concurrency and throughput:** Support for large volumes of concurrent transactions and users without significant performance degradation under high demand
- **Security and compliance:** Encryption, fine-grained access control, auditing, and compliance certifications to protect data and meet regulatory requirements
- **Storage optimization:** Storage architectures tuned for transactional workloads, alongside open table format support and interoperability across engines

- **Multimodel and converged workload support:** Native support for relational, document, key-value, and other data models, and, increasingly, converged transactional and analytical processing and vector data for AI and agentic workloads

## LEARN MORE

### Related research

- *Market Analysis Perspective: Worldwide Database Management Systems Software, 2026* (IDC #US54089426, forthcoming)
- *Market Forecast: Worldwide Database Management Systems Software Forecast, 2026–2030 — Agentic AI, Multicloud Adoption, and Platform Consolidation Drive Market Growth* (IDC #US53749326, September 2026)
- *Market Share: Worldwide Database Management Systems Software Shares, 2025 — Agentic AI and Vendor Consolidation Reshape Competitive Dynamics* (IDC #US54622326, September 2026)
- *IDC Survey: Enterprise Buyer Behavior and Data Management Consolidation Trends* (IDC #US54555726, July 2026)
- IDC's Worldwide Data Management Survey, April 2026 (n = 715)
- *IDC FutureScape: Worldwide Data and Analytics 2026 Predictions* (IDC #US53860225, October 2025)
- *IDC Quick Poll, 2Q25: Agentic AI and Data* (IDC #US53638725, June 2025)

### Synopsis

This IDC study evaluates worldwide operational database software platforms for 2026. Vendors evaluated include AWS, Broadcom, Cockroach Labs, Couchbase, Databricks, EnterpriseDB, Google, IBM, InterSystems, MariaDB, Microsoft, MongoDB, Oracle, Progress Software, SAP, and SingleStore. The market spans relational, nonschematic, and multimodel architectures deployed across on-premises, cloud-native, hybrid, and edge environments, and it is being reshaped by the convergence of transactional and analytical workloads and the rise of agentic AI.

As enterprises move agentic AI into production and consolidate years of accumulated database sprawl, vendors are racing to embed vector search, autonomous operations, and open interoperability as baseline capabilities rather than differentiators. Convergence of transactional and analytical processing is accelerating, deployment is standardizing on hybrid and multicloud, and open source is now nearly universal. The result is a market where transactional reliability is assumed and competitive advantage is shifting to deployment flexibility and real-time intelligence.

"The operational database has become the control point for enterprise AI, the place where agents access and act on live business data," said Devin Pratt, research director of Database Systems for Enterprise Intelligence at IDC. "Leadership will belong to the vendors that make governed, real-time context a core capability, not an add-on."

## ABOUT IDC

International Data Corporation (IDC) is the premier global market intelligence, data, and events provider 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 help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

## Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com

Copyright and trademark notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, and web conference and conference event proceedings. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit idc.com/about/offices. Please contact IDC at customerservice@idc.com for information on additional copies, web rights, or applying the price of this document toward the purchase of an IDC service.

Copyright 2026 IDC. Reproduction is forbidden unless authorized. All rights reserved.