



NUCLEUS  
RESEARCH

# DATABASE TECHNOLOGY VALUE MATRIX 2023

ANALYST

Samuel Hamway

## THE BOTTOM LINE

The database technology landscape is evolving in 2023. Driven by tightened budgets and an increased demand for versatile, efficient solutions, multimodel databases are on the rise. Multimodel databases are particularly popular among SMBs and startups because they consolidate multiple data models into a single platform, simplifying data management with increased scalability. Simultaneously, cloud database vendors continue to stand as a unique segment, providing extensive data ecosystems with cost and integration benefits. In contrast, best-of-breed point solutions continue to offer highly targeted capabilities for specific use cases, underlining the market's diversity. As organizations navigate this complex landscape, they must consider their current needs and future scalability, paying particular attention to their own increasing investments in data-intensive machine learning initiatives.



## OVERVIEW

The database management system market remains diverse in 2023, with no one-size-fits-all solution optimal for all business requirements. Traditionally, the database market has catered to individual business units or small companies managing limited data types with point solutions. These highly specialized, siloed tools have provided data management capabilities for specific, limited use cases, with large enterprises sourcing and connecting

multiple solutions to address data storage needs. However, as macroeconomic headwinds have forced organizations to prioritize costs this year, certain vendors have managed to stand out by offering comprehensive solutions that simplify and streamline data architecture through broad functionality and user-friendly management.

With continued economic uncertainty, particularly in the technology sector, organizations - notably SMBs and startups operating with limited resources - are gravitating towards multimodel databases. The ability to combine multiple data models, such as document, graph, and relational models, in a single data platform offers a competitive edge, enabling scalability, performance, integration, and security advantages over the traditional polyglot persistence paradigm. Conversations with end-users revealed a preference for multimodel systems in scenarios requiring simplified data management, consolidation of data under one roof, and ease-of-scalability considerations for lean teams. This growing focus on multimodel solutions, prompted by market conditions, has led to a noticeable reorientation of the database market.

On the other hand, public cloud vendors constitute a distinct segment within the database market, delivering multiple solutions under a single cost ecosystem. The extensive breadth within their data ecosystems is unparalleled, offering cost consolidation and simplification for data management, as well as native integration between databases and existing solutions within an organization's technology stack. While these ecosystems offer comprehensive capabilities, they also introduce an intricate layer of coordination and an increased IT complexity. The higher the data architecture reconfigurability, the more intricate the system becomes to navigate and manage, subsequently raising the barriers to entry. Cloud vendors that offer native connectivity with a range of commonly used solutions and foster high developer familiarity with their data ecosystems have seen the most significant boosts in terms of usability.

It's important to note that the database market is saturated with best-of-breed point solutions. For vendors not offering consolidation through multimodel capabilities or cost consolidation across a wide range of solutions, their ability to meet specific use cases - whether for a particular data type like graph or relational, or for a specific requirement such as stability or security - can often make point solutions the preferred choice.

In this Value Matrix, vendors are positioned according to the relative usability and functionality of their respective solutions, as well as the value that customers realized from each product's capabilities (Nucleus Research V67 – Understanding the Value Matrix – April 2021) and presented as a snapshot of the current market rather than an empirical ranking of vendors. The arrows indicate perceived momentum in the indicated direction with respect to usability and functionality. Positioning and momentum are informed primarily by conversations with end-users, along with the most recently released capabilities/features and areas of vendor investment.

## LEADERS

Leaders in the Database Technology Value Matrix include AWS, Microsoft, MongoDB, Oracle, and Redis.

### AWS

A Leader in the 2023 Database Technology Value Matrix, AWS supports diverse data models with more than 15 specialized engines. These database solutions include relational, document, graph, key-value, ledger, in-memory, time-series, and wide column databases, provided through services including Amazon RDS, Amazon Aurora, Amazon DocumentDB, Amazon Neptune, Amazon DynamoDB, Amazon Quantum Ledger Database, Amazon ElastiCache, Amazon MemoryDB for Redis, Amazon Timestream, and Amazon Keyspaces. This range of databases allows organizations to select the most cost-effective and scalable solutions for their specific data formats and use cases. AWS databases, often chosen for their availability, reliability, and security, support multi-region deployments, multi-primary replication, and ensure data protection with network isolation and comprehensive encryption. As fully managed services, AWS databases include continuous monitoring, self-repairing storage, and automatic scaling, thus relieving users of tasks like server provisioning, patching, and backups.

Recent product updates and announcements include:

- The Amazon Aurora platform now includes an I/O-Optimized configuration, providing a performance boost and predictable cost structure for data-heavy applications. This update can lead to savings when I/O spend exceeds a quarter of current Aurora database expenditure. Customers can choose between this new configuration and the existing Aurora Standard based on their application's I/O usage. The I/O-Optimized configuration works with Aurora MySQL version 3.03.1 and up, and Aurora PostgreSQL v15.2 and up. Users can switch between configurations every 30 days without needing to reboot their database instances. Currently, this update is available in most AWS regions, with a rollout in China and AWS GovCloud forthcoming.
- Amazon has announced the general availability of Trusted Language Extensions for PostgreSQL (`pg_tle`) on its Amazon Aurora PostgreSQL-Compatible Edition and Amazon RDS for PostgreSQL. This open-source development kit allows developers to safely build high-performance extensions using trusted languages such as JavaScript, Perl, and PL/pgSQL. The new feature provides control to database administrators over who can install and run extensions, and is compatible with Amazon Aurora PostgreSQL-Compatible Edition 14.5 and Amazon RDS for

PostgreSQL 14.5 or later. The project, open for community contributions, is licensed under Apache License 2.0 and is available on GitHub.

- Amazon has launched the general availability of Elastic Clusters for its DocumentDB service, which comes with MongoDB compatibility. Elastic Clusters allows users to scale their document databases elastically, to handle a substantial number of writes and reads and offer petabytes of storage capacity. This new feature simplifies interactions with DocumentDB by automatically managing the infrastructure, eliminating the need for users to manually create, remove, upgrade, or scale instances. Elastic Clusters uses sharding, a database concept that divides large data sets into smaller ones across multiple nodes. This allows databases to scale beyond vertical scaling limits. Elastic Clusters, in addition to offering high scalability, ensures data availability across three Availability Zones (AZs) by default, replicating six copies of the data across these zones. Elastic Clusters is compatible with Amazon DocumentDB API, AWS SDK, AWS CLI, AWS CloudFormation, and the AWS console. It also integrates with other AWS services like Amazon CloudWatch, AWS Identity and Access Management (IAM), Amazon Virtual Private Cloud (VPC), and AWS Glue. Users can create a cluster, scale up or out, and scale in or down using the AWS CLI. Amazon DocumentDB Elastic Clusters is currently available in the following AWS Regions: US East (Ohio, N. Virginia), US West (Oregon), and Europe (Frankfurt, Ireland).
- Amazon has introduced two new features, Optimized Reads and Optimized Writes, for its Relational Database Service (RDS) for MySQL. The Optimized Reads feature aims to expedite query processing by using NVMe-based SSD block storage for temporary tables created by MySQL, leading to query execution being faster. This upgrade is particularly effective for queries that involve sorts, hash aggregations, high-load joins, and Common Table Expressions. Optimized Writes, on the other hand, are designed to double write transaction throughput without additional costs or changing the level of provisioned IOPS, making it a valuable feature for write-intensive workloads. This can include digital payments, financial trading platforms, and online gaming. To use Optimized Reads, customers should select MySQL engine version 8.0.28 or newer and launch Amazon RDS for MySQL on specified instance types. Its performance can be monitored through new CloudWatch metrics. Optimized Reads are available in all AWS regions where the eligible database instance types are available. For Optimized Writes, customers will have to create a new DB Instance from scratch on either a db.r5b or db.r6i instance with the latest version of MySQL 8.0. The feature is currently available in select AWS regions including US East (Ohio, N. Virginia), US West (Oregon), Asia Pacific (Singapore, Tokyo), and Europe (Frankfurt, Ireland, Paris). Both features work to make the RDS

more efficient and cost-effective, continuing Amazon's efforts since its launch in 2009.

- AWS has announced the addition of a fully managed Schema Conversion feature to its Database Migration Service (DMS). This new feature simplifies database migrations by integrating schema assessment and conversion within AWS DMS, eliminating the need for separate downloads and execution of AWS Schema Conversion Tool (SCT). The DMS Schema Conversion function automatically converts source database schemas and most database code objects into a format compatible with the target database. Any objects that can't be automatically converted are clearly marked with instructions for manual migration. Currently, DMS Schema Conversion supports Microsoft SQL Server versions 2008 R2 and higher, and Oracle versions from 10.2 and later, including 11g up to 19c, as sources for migration projects. As for target databases, it supports Amazon RDS for MySQL version 8.x and Amazon RDS for PostgreSQL version 14.x. To use DMS Schema Conversion, users need to create a virtual private cloud (VPC) and set up source and target databases. AWS DMS console will now display options to set up Instance profiles, add Data providers, and create Migration projects. AWS DMS Schema Conversion is now available in select regions including US East, US West, Asia Pacific, and Europe.
- Amazon has introduced a fully managed Blue/Green Deployment feature for Amazon Aurora with MySQL compatibility, Amazon RDS for MySQL, and Amazon RDS for MariaDB. This new feature allows for safer, simpler, and faster database updates by creating a mirrored staging environment that remains in sync with the production environment through logical replication. Blue/Green Deployments allows for significant changes, such as version upgrades or schema modifications, to be tested in the staging environment without impacting the production workload. The staging environment can be promoted to production in as little as a minute with no data loss, thanks to a blocking mechanism that ensures the staging environment is up-to-date before the switchover occurs. The feature can be used via the AWS Management Console, where users can select the database to be updated and initiate the Blue/Green Deployment. Users are charged for the staging environment, including any enabled features like Amazon RDS Performance Insights. Following a successful switchover, the old production environment is not deleted, allowing for further testing if needed, though standard billing charges apply until the old environment is deleted. The Blue/Green Deployment feature is now available in all AWS commercial Regions, excluding China, and AWS GovCloud Regions.

## MICROSOFT

In the 2023 Database Technology Value Matrix, Microsoft has been recognized as a Leader for its diverse database offerings. Azure Cosmos DB, Microsoft's NoSQL database offering, and Azure Cache for Redis, an in-memory database, are also highlighted. Both Azure SQL Database and Azure Cosmos DB support serverless compute, which results in cost savings because customers only pay for the resources they use. Furthermore, Azure Database for PostgreSQL provides scale-out storage and compute, allowing for greater flexibility when scaling servers. Microsoft's range of fully managed database products, which feature automated scalability, availability, and security, enables customers to focus on the health and delivery of their applications while reducing database management time and costs.

Recent product updates and enhancements include:

- Microsoft's Azure Data Manager for Energy, a fully managed service built on the OSDUTM Data Platform, facilitates improved data discovery, security, and scalability for energy companies. In response to user feedback, Microsoft is launching a cost-effective "Developer" tier in mid-June 2023, offering the same security and integration features as the standard version. This tier is suitable for organizations in the initial stages of the OSDU Data Platform or those creating a pre-production testing environment. The Standard tier, recommended for operationalizing workflows or deploying tested models, offers higher database throughput and supports larger production requirements. Azure will support both tiers, with priority given to the Standard tier. Users can access the Developer tier via the Azure Marketplace upon availability.
- Microsoft has announced a preview of Azure Data Manager for Agriculture, an extension of the Microsoft Intelligent Data Platform. The platform aims to address the agriculture industry's mounting challenges by consolidating disparate farm data sources, thereby enabling high-quality datasets for the development of digital solutions. The platform not only fosters innovation but also promotes transparency. This initiative aligns with Microsoft's commitment to sustainability, as exemplified by other projects such as Project FarmVibes and Microsoft Cloud for Sustainability.
- Microsoft has launched a public preview of Azure Managed Lustre, an open-source parallel file system known for high-performance computing (HPC) and cluster computing. This service is designed to aid HPC and AI applications without users having to manage an intricate parallel file system. Azure Managed Lustre offers on-demand cluster deployment serving global regions and provides consistent performance, which resolves the problem of shared infrastructure disruptions often seen on-premises. Users have the option of two durable SSD-based SKUs during the preview phase, providing varying performance choices for crucial workloads. The service also facilitates data tiering between Azure Managed Lustre and Azure Blob,



allowing for data archival, retention, and protection. Azure Managed Lustre is tailored for HPC and AI workloads, offering integration with various Linux distributions and the Azure Kubernetes Service.

- Microsoft's Azure team is partnering with Intel to provide confidential computing on 4th Gen Intel Xeon Scalable processors with Intel Trusted Domain Extensions (Intel TDX), aiming to enable organizations in highly regulated industries to transfer their sensitive data workloads to the cloud. The Intel TDX meets the Confidential Computing Consortium (CCC) standard for hardware-enforced memory protection. Customers across various industries are already using Azure's confidential computing with Intel processors to enhance data privacy and mitigate unauthorized access to sensitive data. They employ solutions like data clean rooms for developing new healthcare therapies and privacy-preserving digital asset management solutions for the financial industry. Azure's current offerings include DCsv3 VMs that provide application isolation using Intel SGX. With the addition of Intel TDX, Azure extends its portfolio to offer isolation at the VM, container, or application levels. Microsoft is also addressing concerns around data sovereignty and regulatory compliance by controlling hypervisors through hardware, which helps protect customer data. In addition, with Intel TDX, attestation is done against the entire VM or container, each possessing a unique hardware key to safeguard memory. Organizations such as Novartis Biome, Fireblocks, and Carbon Asset Solutions have used Azure's existing Intel confidential computing offerings to various beneficial ends.

## MONGODB

MongoDB is a Leader in this year's Database Technology Value Matrix, providing a source-available JSON-like NoSQL database solution. MongoDB employs BSON, a binary format that is a variant of JSON that is more compatible with data types such as date, binary, and augmented number formats. It includes features like real-time embeddable analytics, mainframe offloading, serverless development, and real-time views. MongoDB's adaptable document model can represent a variety of data models, including relational/tabular, key-value, graph nodes and edges, time-series, geospatial, and text, making it a complete data solution. It also includes built-in services for mobile data synchronization, such as Atlas Data Lake, Atlas Search, and MongoDB Realm. MongoDB's data platform is used by clients for a range of applications, including IoT, fraud detection, real-time analytics, customer analytics, and e-commerce. MongoDB's clients come from industries including finance, telecommunications, healthcare, retail, government, payments, manufacturing, and gaming. Clients frequently choose MongoDB because of its dynamic schema with BSON document-based storage, which combines ease of use with JSON-like similarity and high performance with binary optimization. The option for serverless development adds to its appeal by improving scalability and performance.



Recent product updates and announcements include:

- MongoDB, Inc. has announced that its MongoDB Atlas for Government has obtained the formal FedRAMP Moderate Authorized designation. MongoDB Atlas for Government offers tools and services that allow U.S. government organizations to effectively construct and implement secure, large-scale, distributed applications in the cloud. Furthermore, an independent auditor confirmed that MongoDB Atlas for Government is in line with Criminal Justice Information Services (CJIS) requirements. FedRAMP is a U.S. government initiative that establishes a uniform approach for security assessment, approval, and continuous monitoring for cloud products and services. MongoDB Atlas for Government, available since June 2021, is a specialized environment for federal, state, and local agencies, as well as educational institutions. It facilitates the faster development of modern, cloud-native applications while adhering to rigorous data security and privacy standards.
- On October 12, 2022, MongoDB announced it would improve developers' ability to construct data-intensive applications at scale with a pay-as-you-go option for MongoDB Atlas within the Azure Marketplace & Portal. The new offering gives developers access to a free trial and a simplified billing process through their Azure accounts. This development comes as MongoDB participates in the initial set of partners for Microsoft's Intelligent Data Platform Partner Ecosystem, launched at Ignite 2022.
- This collaboration is geared towards supporting developers in the Microsoft ecosystem, providing them with MongoDB Atlas, a platform designed for modern application development. The strategic partnership between Microsoft and MongoDB centers around a shared commitment to a developer-first mindset. This alliance has become increasingly vital in recent years due to the growth in cloud adoption and emerging technologies that allow developers to utilize data more effectively.

## ORACLE

Oracle stands as a Leader in this year's Database Technology Value Matrix, offering a data platform capable of handling diverse data types such as graph, document, key-value, geospatial, temporal, object, and multidimensional data. Oracle Autonomous Database, Oracle's flagship offering, serves as a multimodel, multitenant, and multi-workload database management system, eliminating the need for multiple dedicated databases. Developers can access multimodel data using SQL, REST, and APIs, facilitated by built-in low-code (APEX), developer (SQL Developer), and modeling (SQL Modeler) tools. Oracle also includes Graph Studio, Spatial Studio, AutoML, and ML Notebooks for more accessible development and analysis. The autonomous capabilities automate several operational processes such as

provisioning, configuring, securing, performance tuning, patching, and autoscaling. By automating the addition and removal of indexes for optimal performance, it reduces the manual effort required from database administrators. This fully managed cloud service can be deployed across various public clouds, local clouds, and on-premises in customer data centers, even offering native connection for Azure customers to Oracle Database and Exadata running on Oracle Cloud Infrastructure (OCI). The capabilities of Oracle's platform are further augmented by data security measures, including data encryption at rest and in motion, regulated data protection, security patching, and threat detection.

In addition to the Autonomous Database, Oracle offers the MySQL Database Service, a fully managed service on OCI that automates tasks like backup, recovery, and database and OS patching. It also offers MySQL HeatWave, a query accelerator for MySQL Database Service, enhancing performance for analytics and hybrid workloads. MySQL Autopilot further leverages machine learning to automate data distribution across nodes and optimize workload parallelization, among other tasks. The solution provides administration capabilities to support hybrid cloud environments and simplifies data loading, transformation, and insights. Oracle's suite is also complemented by consulting services to assist with deployments.

Recent product updates and enhancements include:

- Oracle released the Oracle Database Service for Microsoft Azure in July 2022, allowing Azure customers to easily provision, access, and manage Oracle Database services and Exadata in OCI via Azure. This builds on previous Oracle innovations. With Azure Interconnect, customers with multi-cloud strategies involving Oracle and Microsoft Azure can further reduce costs and complexity. The service configures networking automatically and allows Azure users to use Azure security to access Oracle Database and Exadata in OCI without incurring data egress charges.
- Oracle's Autonomous Database now includes a new built-in application called Data Studio. This unified interface consolidates the platform's data tools, enhancing data loading, understanding, transformation, and analysis, along with providing immediate data insights and secure data sharing options. Users can leverage these features either independently or alongside their preferred tools. For instance, Analytic Views prepared using the Data Analysis tool can be utilized to optimize analytic queries in PowerBI or Tableau. Data Studio also allows users to construct Extract-Load-Transform (E-L-T) data flows with the Data Transforms feature, and share data securely using the open Delta Sharing standard with Data Share. Looking ahead, Oracle plans to expand Data Studio's capabilities, incorporating user interfaces for the Graph, Spatial, and Machine Learning functionalities of the data platform, aiming to maximize user productivity without the need for additional service setup.

- Oracle has released Zero Downtime Migration (ZDM) 21.4, an update to its solution for moving Oracle Database workloads to Oracle Cloud. This version introduces enhancements in physical migration capabilities, sudo-less migration, and migration of objects with different GoldenGate support modes. Physical Migration enhancements include Redo Apply Catch Up for Paused Migration Jobs, enabling ZDM to monitor the redo apply lag after the standby database's creation. ZDM now acknowledges user-initiated manual Data Guard switchovers and resumes the automated workflow. The Recovery Manager section size for backups can be configured manually, and a new post-migration task allows the upgrading of the target database time zone file. Logical Migration improvements include separate phases for data and metadata migration, offering a new workflow for user and profile creation, metadata creation, and data import. ZDM 21.4 now supports tables with XML data types and allows migration via Data Pump's automated conversion to Binary Storage. A new authentication plugin facilitates sudo-less migration, enabling databases to be migrated without non-privileged users for both source and target node access. Other updates encompass enhanced Data Pump error handling, Object Storage Pre-authenticated URL support, and auto-start and restart profiles for Oracle GoldenGate microservices, among other features.
- Oracle Cloud Infrastructure (OCI) has announced the general availability of its Full Stack Disaster Recovery service. This is OCI's first true disaster recovery-as-a-service (DRaaS) solution, enabling comprehensive disaster recovery management for an entire application stack. The service manages the transition of infrastructure, platforms, and applications between OCI regions globally, allowing the deployment of disaster recovery environments without having to redesign or redeploy existing infrastructure, databases, or applications. Full Stack Disaster Recovery can be accessed via the Oracle Cloud Console under Migration & Disaster Recovery as "DR Protection Groups". The service revolves around the concept of disaster recovery protection groups (DRPG). Currently, Full Stack Disaster Recovery is available in Amsterdam, Ashburn, Frankfurt, London, and Phoenix regions, with plans to expand the service to all realms and regions within the next six to nine months. The service is flexible, extensible, and scalable, ideal for cross-regional disaster recovery in OCI, including moving workloads between OCI regions for maintenance, rolling upgrades, and other uses. The service automatically creates basic disaster recovery plans based on the virtual machines, storage volume groups, and databases included in a DRPG, significantly reducing configuration effort and time. Additional resources such as the User Guide and Live Lab are available to further support users.

## REDIS

Redis is a Leader in this year's Value Matrix. Redis is an open-source, data platform supporting an array of abstract data types such as strings, lists, hashes, sets, sorted sets, HyperLogLogs, bitmaps, streams, and spatial indices, among others, enabling direct operations on these data types. Additionally, Redis includes a set of capabilities including a publish-subscribe messaging system, clustering, and master-replica replication. Redis Inc's products are further augmented by their compatibility with major cloud service providers like Microsoft Azure, Google Cloud, and AWS, which facilitates support for multi-cloud, hybrid cloud, and microservices infrastructures.

Redis Enterprise, the commercial version offered by Redis Inc., is a multimodel data platform offering a distributed system that is often deployed to support high-speed transactions, fraud detection, recommendation engines, ML, IoT, search, real-time analytics, and e-commerce applications.

Redis Enterprise's search functionality is equipped with capabilities for secondary indexes, a query engine, and full-text search capabilities. In addition to this, Redis Enterprise presents a range of other capabilities including Triggers and Functions that provide developers a way to process their data within Redis bringing the business logic closer to the data and taking advantage of advanced data structures and algorithms such as Bloom filters, Cuckoo filters, Count-Min Sketch, and Top-K. All these features contribute to the platform's flexibility and robustness.

The platform supports multiple built-in data structures, document, vector, and additional data models, enabling the development of modern applications and the generation of insights across different modes of data. Furthermore, Redis Enterprise supports languages such as Python, Java, C, and JavaScript, allowing programming any functionality across data models.

Recent product updates and announcements include:

- The vendor introduced enhanced JSON support, a feature that increases flexibility in handling document databases. The enhanced support allows for more diverse and intricate data storage and retrieval, leading to optimized efficiency in customer data operations.
- Redis recently unveiled an update which includes native indexing, querying, and full-text search features. By simplifying complex tasks for developers, these new features enable a swifter and more efficient development process for applications.
- Redis has added new functionality to Search patterns, including multi-value indexing. This addition provides a significant improvement in data retrieval capabilities, potentially leading to more precise and valuable business insights.

- Redis Enterprise now offers enhanced querying, secondary indexing, and full-text search capabilities beyond JSON. These capabilities enable customers to explore and retrieve data from a wider variety of data types and sources, thereby improving search efficiency.
- The vendor has improved support for AI workloads, introducing native support for Vector Similarity Search (VSS). This is particularly beneficial for applications that employ large language models, such as recommendation engines, consequently enhancing user experience and providing more potent business intelligence.
- Redis has made Active-Active deployment on Kubernetes easier with the introduction of the Redis Enterprise Operator. This tool simplifies the management of complex architectures, significantly reducing administrative workload.
- Redis introduced version 7.0 in the fully managed Redis Enterprise Cloud service. The update brings enhanced server-side processing functionality and greater control of Pub/Sub channels, offering significant improvements in security and efficiency.
- The partnership between Redis and Microsoft has resulted in the Azure Cache for Redis Enterprise service. This collaboration ensures higher availability and business continuity, providing a reliable and resilient service for customers.
- Redis recognized the importance of developer experience and enhanced it by acquiring RESP.app. This improvement aims to speed up development and improve efficiency, potentially leading to faster delivery of customer-facing applications.
- Redis has extended its strategic partnership with AWS. This expanded collaboration is designed to accelerate cloud migration and application modernization, providing potential cost savings and performance improvements for customers.
- Redis Stack has been updated to increase developer productivity with Client Libraries. This comprehensive package aims to provide a complete developer experience, accelerating application creation and improving debugging.
- Redis has increased the flexibility of deploying the Redis Enterprise Cloud by introducing Pulumi support. This tool allows developers and administrators to manage cloud resources more effectively, ensuring consistency and efficiency across different environments.

## EXPERTS

Experts in the Database Technology Value Matrix include Alibaba Cloud, Google, IBM, InterSystems, MarkLogic, and SAP.

## ALIBABA CLOUD

Alibaba Cloud, recognized as an Expert in this year's Database Technology Value Matrix, provides a comprehensive suite of database products. At the core of the suite is PolarDB, a cloud-native relational database that supports MySQL, PostgreSQL, and Oracle. This platform combines the performance and reliability of traditional enterprise databases with the flexibility of open-source databases. Alibaba Cloud offers ApsaraDB RDS for MySQL, PostgreSQL, SQL Server, and MariaDB. The vendor also provides PolarDB for Xscale, a service tailored to handle transactional business data. Additionally, Alibaba Cloud has developed Lindorm, a cloud-native database service with elasticity and multimodel data support. For high concurrency, large data storage, and performance bottlenecks, Alibaba Cloud offers PolarDB-X. The vendor also maintains Tair, a Redis-compatible in-memory database service, providing features such as Global Distributed Cache, data flashback, and Transparent Data Encryption (TDE). To facilitate database management, Alibaba Cloud includes the Data Transmission Service (DTS) for migration and synchronization tasks, Database Backup (DBS) for data protection, Data Management (DMS) for data security and collaboration, and the Database Autonomy Service (DAS) for self-optimization and self-security measures.

Recent product updates and enhancements include:

- Alibaba Cloud announced an update to PolarDB for MySQL, incorporating a serverless function for non-serverless clusters. This new functionality is aimed at improving resource elasticity. It allows the quick and automatic addition of serverless elastic resources when the current system is overloaded, enhancing throughput performance and system stability to handle traffic surges, without the need for auto-scaling configuration.
- The vendor also announced the release of the RDS for MySQL Cluster Edition. This edition includes additional features like readable standby nodes and multi-source replication (MGR). These enhancements are especially beneficial for scenarios that demand high read volumes but have minimal write operations, providing cost efficiency.
- Alibaba Cloud released RDS for PostgreSQL Serverless, designed to offer enhanced flexibility and cost-effectiveness.
- Alibaba Cloud announced that ApsaraDB for MongoDB will now support the modification of instance service ports. This update aims to streamline cloud migration processes by removing the need for application reconfiguration when the default MongoDB port isn't in use in the production environment.

## GOOGLE

In the 2023 Database Technology Value Matrix, Google Cloud is recognized as an Expert. It provides a variety of database services including relational options like Cloud SQL, AlloyDB for PostgreSQL, and Cloud Spanner, key-value storage with Cloud BigTable, document storage via Firestore and Firebase, and in-memory storage with Memorystore. These databases cater to a wide range of sectors such as retail, manufacturing, automotive, energy, healthcare, media, telecommunications, government, public sector, and financial services. Google's databases, are easily scalable, and have native integrations with other Google Cloud services, including BigQuery, Google Kubernetes Engine (GKE), and the Apache ecosystem. Offering encryption for data at rest and in transit, these databases ensure high security. Their high availability ensures minimal disruption to applications and processes. Cloud Spanner, Cloud Bigtable, and Firestore, for example, offers up to 99.999% availability SLA.

Recent product updates and enhancements include:

- Google has recently previewed AlloyDB Omni, a downloadable edition of AlloyDB that can function on-premises, at the edge, across different clouds, or on developer laptops. This product is driven by the same engine as the AlloyDB service, a fully managed, PostgreSQL-compatible database. According to Google, AlloyDB Omni is twice as fast as standard PostgreSQL for transactional workloads and delivers more rapid analytical queries. It leverages AlloyDB's columnar engine to minimize latency, allowing for speedier data scanning, joining, and aggregating. The system also incorporates machine learning for efficient data organization. Moreover, AlloyDB Omni features an index advisor that optimizes frequently run queries, reducing the tuning guesswork for query performance. As a PostgreSQL replacement, it supports all compatible tools for database backup and replication. It also aligns with any PostgreSQL-enabled applications businesses are currently using. AlloyDB Omni will be generally available later this year and will offer full enterprise support including 24/7 technical support and software updates.

## IBM

Recognized as an Expert in the Database Technology Value Matrix for 2023, IBM's database offerings are suitable for organizations looking to extend their relational database to support multimodel data in a hybrid and multicloud environment. The company offers various fully-managed and highly-available cloud databases, such as IBM Cloud Databases for PostgreSQL, IBM Cloud Databases for MySQL, and IBM Db2 on Cloud. These offerings provide autoscaling for PostgreSQL and MySQL, serverless capability for Db2, and multiple availability zones for lower latency processing and minimized local outage risks. They also come with advanced security features, including data encryption at rest and in transit. IBM's



flagship Db2 is a versatile data platform supporting multiple data models such as XML, JSON, BSON, graph, and relational, accessible via SQL and APIs. Db2 is equipped with AI capabilities facilitating no-code model creation and training. It is commonly used for a range of applications like IoT, mobile, big data, fraud detection, real-time analytics, and operational reporting. IBM's vision emphasizes an integrated multimodel Db2 platform, featuring AI, self-service, and automation capabilities.

Recent product updates and announcements include:

- IBM has introduced a Db2 operator for Kubernetes on the AWS Elastic Kubernetes Service (EKS). This move marks an increase in cloud-centric strategy for the database; Db2's transition to the cloud has been a gradual process, beginning with the launch of a cloud-native version of the database, Db2u, which was initially only accessible on RedHat OpenShift. However, IBM has now expanded Db2u universal container services to Amazon EKS, with expansion to other clouds expected.
- IBM has acquired StepZen Inc, a company known for creating a unique GraphQL server, which simplifies the building of GraphQL APIs. This acquisition complements IBM's ongoing initiatives in database management, data fabric, API management, and integration. The uptake of hybrid cloud solutions has dispersed data across various locations, highlighting the importance of APIs as bridges between disparate data sources and applications. As GraphQL emerges as a popular query language for interacting with data, this acquisition will strengthen IBM's position in the database market by enhancing the accessibility and interactivity of its data solutions.

## INTERSYSTEMS

InterSystems is recognized as an Expert in this year's Database Technology Value Matrix. InterSystems' IRIS is a cloud-first data platform that supports a variety of data models such as relational, object, document, key-value, cubes, and multidimensional arrays. InterSystems IRIS combines several technologies into one platform, which simplifies architecture, reduces the need for coding, lowers resource usage, and minimizes maintenance efforts. It is used in different applications including customer analytics, ML apps, risk analytics, IoT, and specialized applications like healthcare. IRIS builds upon the functionalities of InterSystems Caché and Ensemble and introduces new features to develop and deploy analytics-focused enterprise applications. The platform provides multiple APIs for transactional persistent data operations and supports a range of languages including SQL, Java, node.js, .NET, C++, Python, and server-side ObjectScript language. InterSystems IRIS also includes features for horizontal scalability, such as sharding and ECP, high availability, business intelligence, transaction support, and backup.

Recent product updates and enhancements include:

- The vendor provided an Analytics and ML enhancement with the introduction of Columnar Storage, a new storage option for InterSystems IRIS SQL tables. Columnar Storage offers a flexible option for running analytical queries much faster than with traditional row storage on InterSystems IRIS. It is fully supported for production use in 2023.1, excluding columnar storage for sharded tables.
- InterSystems announced a new capability for leveraging external data in InterSystems IRIS. The Foreign Tables feature presents itself as regular InterSystems IRIS tables to any SQL queries you write, though the data is not physically stored within the InterSystems IRIS server. It is currently available as an experimental feature and is not yet supported for production environments.
- The vendor released updates for analytics capabilities with InterSystems IRIS, notably InterSystems Reports, Adaptive Analytics, and IntegratedML.
- InterSystems released scalability enhancements that enable large production deployments to meet highly demanding workloads. These include the asynchronous reading of journal files during de-journaling and changes to the Enterprise Cache Protocol (ECP) infrastructure to optimize resource usage and limit contention under high load.
- InterSystems announced that the 2023.1 release adds support for macOS 13 (Ventura) as a new server platform.

## MARKLOGIC

MarkLogic is recognized as an Expert in this year's Database Technology Value Matrix. MarkLogic is a multimodel data platform with a range of capabilities, providing support for various data models like XML, JSON, text, geospatial, and RDF. This platform can manage a spectrum of use cases, including IoT, big data, fraud detection, real-time analytics, edge computing, and content management. The platform offers enterprise-grade scalability, analytics, search capabilities, data security, and reliability. It's designed for hybrid-cloud services, self-automation, AI/ML, and distributed data hubs. The MarkLogic Data Hub Platform is capable of integrating enterprise data to provide business value, running either on-premises or in public or private cloud environments. MarkLogic's features include advanced indexing capabilities, supporting tokenization, collation, and stemming for over 200 languages. It also supports full-text search across data and metadata using various languages like JavaScript, XQuery, SPARQL, and SQL. Other features include RDF triple usage for semantics, ACID compliance, replication for high availability, sharding for horizontal scalability, and built-in security features such as element-level permissions and data redaction. Moreover, its Optic API allows developers to view their data as documents,

graphs, or rows. MarkLogic can maintain a flexible data model with its distributed, scale-out architecture that can handle large amounts of documents and data.

Recent product updates and enhancements include:

- The vendor announced that the Optic API, initially developed for MarkLogic 9, has been upgraded in MarkLogic 11. This API, which was designed specifically for querying multimodel data, now supports geospatial analysis as well. It enables comprehensive indexing and querying of geospatial data and ensures compatibility with GIS tools, including support for OpenGIS and GeoSPARQL.
- MarkLogic announced its compatibility with various BI tools through its GraphQL connectivity. This functionality facilitates the exposure of data modeled and queried via the Optic API in GraphQL, simplifying developer workflows and augmenting the platform's interoperability with existing tools.
- The vendor announced solutions for managing the challenges posed by growing data volumes with MarkLogic 11. The Admin UI has been revamped to improve system observability, auditability, and manageability. Moreover, adaptive memory algorithms have been introduced to handle larger data query result sets effectively. The update also introduces cold storage detection to maintain optimal performance despite potential latencies.
- MarkLogic enhanced security features in version 11. Enhancements include improved backup processes, advanced certificate management for intra-cluster communication, and OAuth support for managing encryption keys. The updated security features complement the existing enterprise-level features of the platform, making it a trusted choice for data management.
- MarkLogic announced improved cloud strategy support with the release of MarkLogic 11. This includes availability as cloud formation templates and Docker images. The update also introduces support for Kubernetes, offering an efficient solution for deployment and orchestration. The company plans to release further enhancements in this area in the upcoming year.

## FACILITATORS

Facilitators in the Database Technology Value Matrix include Cockroach Labs, Couchbase, MariaDB, Neo4j, and Singlestore.

## COCKROACH LABS

Cockroach Labs' CockroachDB is recognized as a Facilitator in this year's Database Technology Value Matrix for its PostgreSQL-compatible distributed database.

CockroachDB, developed and maintained by Cockroach Labs, serves a wide range of industries, including finance, gambling, gaming, manufacturing and logistics, retail and eCommerce, SaaS, and startups. 3x data replication, automatic rebalancing, hotspot detection, automatic data distribution, row-level tuning, spatial data types, nativized JSON, geographic data regulation compliance, and custom availability goals are among the notable features. Furthermore, the software can be deployed on multi-cloud infrastructure orchestrated with Kubernetes. Customers choose CockroachDB for its library of security and downtime-resistant features in data distribution, replication, and archiving to ensure resilience for petabyte-scale configurations.

Recent product updates and announcements include:

- Cockroach Labs has introduced a new database migration tool called Molt to their SQL database. Molt aims to simplify the database migration process by mitigating associated complexities. It offers a new schema conversion tool that identifies and rectifies incompatibilities between the source database and CockroachDB. The process of database migration is often resource-intensive and challenging, with enterprises frequently encountering issues such as lack of interoperability and incompatibility, which can lead to data inconsistencies. Molt is designed to alleviate these issues, thereby streamlining the migration process.
- Cockroach Labs has made CockroachDB Serverless generally available. This on-demand relational database uses a PostgreSQL interface, offers consumption-based elastic scaling and pricing, and includes features aimed at decreasing database operations and budget expenditure. Additionally, it provides developer tools like a command-line interface (CLI) integrated with PostgreSQL object-relational mapper (ORM) software. ORM is a code library that enables developers to query and manipulate database data using an object-oriented paradigm, thus permitting communication with the database in the same language without relying on SQL. Other serverless databases include Azure SQL Database, Amazon Aurora Serverless, and Planetscale, with Oracle Database on OCI (Oracle Cloud Infrastructure) also offering a similar function despite not typically fitting the serverless pattern.
- Cockroach Labs has announced integrations with HashiCorp, Vercel, and Terraform products. The integration with Vercel, currently in private preview, aims to assist developers in creating and deploying web applications with a relational database, eliminating the need for management or maintenance. The HashiCorp integration will help developers manage database credentials and encryption keys using HashiCorp Vault's Dynamic Secrets. The integration with Terraform, also in private

preview, aims to simplify the provisioning of CockroachDB Dedicated & Serverless in the cloud via Terraform Provider, thus facilitating DevOps teams in deploying applications.

## COUCHBASE

Couchbase, a Facilitator in this year's Database Technology Value Matrix, provides a multimodel NoSQL database used across several industries such as gaming, healthcare, and retail. The database supports an array of data models, including key-value, document, text, search, and spatial. It employs a SQL-like query language, offering a familiar environment for relational database developers, and offers software development kits (SDKs) for various languages, facilitating application development. Couchbase can be deployed as a Database-as-a-Service (DBaaS) through the vendor's Capella offering. Couchbase can be deployed on both public and private clouds, as well as containers and bare metal servers. Its functionalities range from full-text search to distributed transactions and multi-cloud data management. The platform is primarily used for customer profiling, social networking, Internet of Things (IoT), mobile, and operational workloads.

Recent product updates and enhancements include:

- Couchbase announced updates to Capella; these updates aim to streamline workflows and increase productivity by integrating with popular developer platforms, such as Netlify and Visual Studio Code (VS Code). Furthermore, the enhancements include support for time series data, expanding usability for a variety of applications like IoT or finance apps. Capella's deployability has been broadened across major cloud service providers with added security and compliance capabilities. Other new management features include dynamic disk scaling, hibernation of clusters, and the provision of downloadable buckets.
- Couchbase has announced the integration of Couchbase Capella with the Microsoft Azure Marketplace. This new availability enables streamlined deployment and management of Capella on the Azure cloud platform, providing customers with the flexibility to leverage Azure's cloud infrastructure.
- Couchbase has introduced the Independent Software Vendor (ISV) Starter Factory, a program designed to equip ISVs with the necessary tools and resources to develop and monetize their applications using Couchbase Capella on Amazon Web Services (AWS). The program, a collaboration between Couchbase, AWS, and selected System Integrators (SIs) from the AWS Partner Network, aims to streamline ISV migration to Capella on AWS. This initiative builds on the recent enhancements to Couchbase's ISV Partner Program and the expanded alliance with AWS, offering

customers integrated marketing activities, commercial incentives, and technology integrations, including support for transitioning workloads to Capella on AWS.

- Couchbase has updated its Independent Software Vendor (ISV) Partner Program, introducing a new Success Package and ISV Program Guide. The Success Package includes improved training, certifications, and resources to aid ISVs in developing and monetizing their applications. The Couchbase Capella cloud data platform, a core feature of the program, allows ISVs to consolidate their architecture, expedite application development, and decrease Total Cost of Ownership (TCO). It facilitates the integration of Couchbase's cloud database platform with the ISVs' applications or other solutions. The program supports ISVs in reducing data duplication and accelerating release cycles through the use of Capella's JSON document and key-value database.

## MARIADB

MariaDB Corporation Ab, recognized as a Facilitator in this year's Database Technology Value Matrix, offers the fully managed cloud database service MariaDB SkySQL. This service executes MariaDB as a distributed SQL database using the MariaDB Xpand platform, providing functionalities such as transactional, analytical, and hybrid workload support, dynamic SQL scaling, ACID-level consistency, and top-tier reliability. Deployable on both Google Cloud and AWS, SkySQL facilitates multi-cloud and hybrid-cloud setups, often selected for its MariaDB functionality support in production deployments. Originating as a community-developed, commercially supported offshoot of MySQL, MariaDB offers compatibility with MySQL and supports relational, columnar, geospatial, JSON, and graph data structures. The platform can interface with several programming languages like C/C++, Java, and .NET. Users typically employ MariaDB for varied workloads, including comprehensive customer analytics and operational reporting.

Recent product updates and enhancements include:

- Addressing the limitations of traditional geospatial solutions, MariaDB introduces Geospatial, an API-driven, developer-friendly platform based on an acquisition made last year. MariaDB Geospatial employs open standard interfaces for web mapping and geospatial queries, fusing vector storage in the database with raster storage in cloud object storage, creating a flexible solution that doesn't necessitate a definitive choice. MariaDB Geospatial is currently available for testing on SkySQL. Customers see geospatial data processing as increasingly crucial due to the escalating volume of data from sources like drones, satellites, and sensors.
- In addition to running databases as a service, SkySQL now also offers automatic backup for these databases, a feature typically found in other cloud databases. In

line with its recent observability enhancements, MariaDB is extending this cloud backup service to on-premise databases and databases in other cloud infrastructures. This extension allows databases outside the cloud to harness cloud benefits, such as disaster recovery strategies with the option to restore backups in the user's preferred cloud. This Cloud Backup service for remote databases is now available for testing on SkySQL.

- MariaDB has extended its SkySQL Observability service to include MariaDB Community Server deployments, allowing a unified view of all MariaDB databases deployed across different environments. This enhancement aids in rapid problem detection and analysis. SkySQL Observability offers deep insights into MariaDB products and supports strategies such as disaster recovery failovers. The service features real-time performance dashboards, log collection, event management, and alert notifications. It also integrates with third-party services like Datadog, AppDynamics, and New Relic through an observability API. It is currently available and free for databases under a MariaDB subscription. For additional databases, including MariaDB Community Server ones, pricing begins at \$45 per month per database node.
- MariaDB has introduced Serverless Analytics, a feature designed to extract insights from current data with a pay-per-use model, eliminating the need for Extract, Transform, Load (ETL) processes. Alongside this, an Autoscaling feature has been enabled, dynamically scaling resources up or down based on demand fluctuations, optimizing costs. When combined with MariaDB's distributed SQL database, Xpand, autoscaling effectively manages unexpected demand surges.

## NEO4J

Neo4j, a Facilitator in this year's Database Technology Value Matrix, is a graph database management system produced by Neo4j, Inc. This ACID-compliant transactional database natively stores and processes nodes, edges, and their attributes, providing a platform for understanding connections and patterns within data. Neo4j's graph database is designed to represent connections between data points, aligning with the mathematical structure of graphs. Available in a community edition under a modified GNU General Public License, it also has an expanded version under a closed-source commercial license that includes online backup and high availability extensions. Neo4j's advanced offerings encompass a variety of functionalities such as multi-data center support, comprehensive security features, graph analytics, and visual graph discovery. These features allow customers to leverage the solution for use-cases like real-time recommendations, AI, graph-based search, data science, and master data management. Neo4j's AuraDB, the vendor's fully managed solution, offers global availability and multi-cloud freedom. With compatibility across AWS,



Microsoft Azure, and Google Cloud, AuraDB can be deployed in over 60 global cloud regions. Its architecture ensures native migration across clouds, and consistent operation worldwide, thereby preventing cloud vendor lock-in and enhancing users' flexibility of choice.

Recent product updates and enhancements include:

- Neo4j recently made a strategic move to acquire 11 full-time specialists from Distributed Technology Associates (DTA), marking a significant expansion of its global Cloud-Managed Services (CMS) capabilities. This move enables Neo4j to offer its enterprise customers enhanced deployment models and accelerated execution, aligning with the rapid scaling of graph technology across hybrid and multi-cloud environments.
- Neo4j announced the launch of Neo4j 5, a next-generation cloud-ready graph database. It promises increased performance and improved scalability across all deployment types - on-premises, cloud, hybrid, or multi-cloud. The update aims to enable rapid development and large-scale deployment of intelligent applications, aiming to maximize data value.
- Neo4j announced the attainment of Amazon Web Services (AWS) Data and Analytics Competency status. This validates Neo4j's success in aiding customers with effective data collection, storage, governance, and analysis at any scale. The status distinguishes Neo4j as an AWS Partner Network (APN) partner with specialized software for complex AWS-based projects, affirming its deep expertise in AWS and its ability to deliver seamless solutions on this platform.

## SINGLESTORE

SingleStore is included as a Facilitator in this year's Database Technology Value Matrix, recognized for its distributed relational database system that's compatible with Apache Spark, Apache Kafka, Microsoft Azure, Google Cloud, and AWS. This platform has a presence in the financial, media, energy, telecommunications, and technology sectors. SingleStore combines operational analytics, machine learning, and transactional capabilities in a single engine, allowing users to drive rapid, interactive analytics and machine learning for their applications without any data movement. Users frequently select SingleStore's distributed database for managing large data sources to power applications and real-time analytics.

With the SingleStore Pipelines, the platform can ingest streaming data with live deduplication, applying parallel processing from different data sources, which may lessen the need for ETL middleware. With its three-tier architecture, SingleStore offers separation of storage and compute and unlimited storage by being able to move data between its in-

memory row-store, disk-based column-store and the cloud object store. SingleStore has been offering vector capabilities as part of its database since 2017, and helps customers deliver on machine learning use cases including semantic search, image matching, facial recognition, and more. SingleStore also applies techniques such as lock-free skip lists, column store compression, and low-level code query compilation to optimize performance. It includes security features such as encryption, audit trail, and security configurations at a petabyte scale, making it a fit for enterprise deployments requiring high security.

Recent product updates and enhancements include:

- On June 6, 2023, The 8.1 version of SingleStore was released. A key addition includes SingleStore Kai for MongoDB, an API that delivers real-time analytics for MongoDB applications. It also improved on its multimodel capabilities for handling semi-structured and unstructured data. For cloud-native enhancements, it introduced a native tool for data ingestion from MongoDB collections and bolstered security with OpenSSL 3.0 support. Observability was also improved with the introduction of a Grafana dashboard and a native alerting feature. The distributed SQL capabilities now include a new type of index called Column Groups for faster data retrieval. To enhance developer experience, an AI assistant named SQrL was introduced, along with an extension to Jupyter Notebooks and a new driver for the SQLTools VSCode extension. The partner ecosystem was strengthened with new collaborations with Oracle GoldenGate, Tableau, and Qlik.
- On May 25, 2023, SingleStore introduced KAI, a MongoDB API. This solution addresses MongoDB's limitations in delivering effective analytics on its JSON/BSON data. KAI helps to decrease database sprawl and complexity, allowing companies to utilize MongoDB data in real-time and conduct real-time analytics without any changes to MongoDB queries.
- In Summer 2022, During the (r)evolution 2022 event, SingleStore announced several new features, including Unlimited Storage (Bottomless), SingleStore Workspaces, and the SingleStore Code Engine Powered by Wasm.

## CORE PROVIDERS

Core Providers in the Database Technology Value Matrix include ArangoDB, EnterpriseDB, ScyllaDB, and TigerGraph.

## ARANGODB

ArangoDB is a Core Provider in this year's Database Technology Value Matrix. ArangoDB is an open-source native multimodel platform that supports multiple data models including document (JSON), graph, and key-value, all unified under a single declarative query language - ArangoDB Query Language (AQL). ArangoDB's multimodel capability enables it to address a wide range of use cases, including microservices, fraud detection, transactional workloads, and data services. The database system supports on-premises, public cloud, multicloud, private cloud, and hybrid cloud deployments. Though ArangoDB is a NoSQL database system, its query language AQL is similar to SQL. This allows a combination of different data access patterns in a single query. ArangoDB's ArangoGraph Insights Platform, previously known as Oasis, provides a DBaaS solution, offering fully-fledged ArangoDB cluster deployments without the need for users to manage or operate the system independently. The platform allows users to run their databases on their preferred cloud provider's data centers, including Google Cloud Platform (GCP), Amazon Web Services (AWS), or Microsoft Azure. This arrangement enhances performance while reducing costs. The ArangoGraph Insights Platform enables the constant availability and health of users' databases through 24/7 monitoring. Additionally, it facilitates the automatic installation of new updates, thereby ensuring databases remain current without service interruption. In terms of data security, the platform offers encryption, audit logs, and frequent data backups. Moreover, the platform guarantees that data ownership remains with the user, and safeguards access using industry-standard protective measures.

Recent product updates and enhancements include:

- ArangoDB announced its achievement of SOC 2 compliance, demonstrating its commitment to stringent data security and availability protocols. This ensures the secure management and operation of its client's database services, allowing them to focus on their core business while ensuring that their valuable data assets are protected. The compliance verifies ArangoDB's adherence to recognized best practices in data security.
- The latest release, ArangoDB 3.10, introduces several significant features and improvements to enhance its database capabilities. It now supports the native ARM architecture, including compatibility with Apple silicon, such as the M1 chips, extending beyond its existing x86-64 architecture support. This update also enables users to test SmartGraphs and SatelliteGraphs on a single server, then transfer them to a cluster, simplifying the testing and deployment process. The release unveils EnterpriseGraphs, a specialized version of SmartGraphs that offers an automated sharding key selection feature for more efficient graph creation. Furthermore, the renaming of (Disjoint) Hybrid SmartGraphs to (Disjoint) SmartGraphs using SatelliteCollections was conducted without altering their functionality. Finally, the

addition of the 'Computed Values' feature allows users to define expressions at the collection level, which can automatically run during data insertions or modifications. This offers more flexibility in handling data attributes and indexing. In addition to these new features, ArangoDB 3.10 also incorporates numerous bug fixes, increasing the platform's stability.

## ENTERPRISEDB

EnterpriseDB, recognized as a Core Provider in this year's Database Technology Value Matrix, provides its DBaaS solution, BigAnimal, which enhances PostgreSQL's capabilities for handling enterprise-level workloads. This makes EnterpriseDB a popular choice among customers for managing large-scale operations due to its advanced tools and PostgreSQL familiarity. BigAnimal offers features such as portable licensing, compatibility with AWS and Azure, Oracle SQL-compatible functions and syntax, continuous backups, role-based access control, activity logs, and connectivity to Apache Superset and pgAdmin. EDB Postgres Advanced, EnterpriseDB's proprietary distribution of PostgreSQL, integrates enhancements in performance, security, and manageability, accommodating enterprise-class workloads. It features unique Oracle compatibility, facilitating easy migration of Oracle databases with minimal modifications. PostgreSQL supports various data types, including relational, key/value, XML, geospatial, and document. Customers utilize EDB for diverse use cases, including e-commerce, operational, and transactional workloads.

Recent product updates and enhancements include:

- On June 20, 2023, EnterpriseDB announced updates to its BigAnimal platform, including rapid deployment of fully-managed Postgres clusters on any primary cloud platform, either via customer or BigAnimal's cloud accounts. This upgrade empowers businesses with options for Postgres distribution and Oracle compatibility, offering versatility without vendor lock-in. Customers can now easily establish and deploy Postgres clusters on any main cloud platform, offering increased data control and flexibility. BigAnimal's new features include high availability with distributed Postgres, promoting optimal uptime for critical applications. It ensures migration ease with Oracle compatibility, allowing swift migration to Postgres from legacy applications.
- EnterpriseDB announced on April 18 that it is offering free on-demand training for Postgres. This initiative aims to address the lack of skilled Postgres professionals. The training resources cover various aspects such as database administration, performance optimization, SQL programming, and application development, catering to different skill levels. The courses, developed by industry experts, are delivered through online platforms and interactive workshops. EnterpriseDB's free training initiative is designed to help individuals and businesses maximize their use

of the open-source Postgres database. The resources are available on the EDB Training Portal.

- EnterpriseDB has introduced the EDB Guaranteed Postgres Migration program, which aims to facilitate the migration process for organizations transitioning from Oracle databases to Postgres. Postgres, recognized for its flexibility and compatibility with various cloud platforms, has become a popular choice for enterprises seeking cloud-based database solutions. The migration program from EnterpriseDB offers a streamlined process to shift from Oracle-based applications to Postgres. This allows organizations to both develop new cloud-native applications and transition existing ones to Postgres, potentially resulting in cost savings.
- EnterpriseDB has launched EDB Postgres Distributed (PGD) 5.0, a product designed to decrease both planned and unexpected downtime for crucial business applications. The new version aims to provide up to very high availability for PostgreSQL databases, using active-active technology, an upgrade that isn't offered by other Postgres-native databases.
- EnterpriseDB has made Transparent Data Encryption (TDE) available globally. This feature enhances Postgres databases with robust security and performance capabilities. TDE improves data security and reduces risks, making it particularly attractive to large enterprises.
- EnterpriseDB has launched its EDB Tools and Extensions Release for PostgreSQL 15, aimed at facilitating the use of Postgres as a standard enterprise database. The release is designed to aid enterprises in developing modern applications utilizing PostgreSQL 15.1. The flexibility of EDB PG 15 extends to various deployment environments, including on-premises, in the cloud, and self-managed or fully managed through EDB BigAnimal, EnterpriseDB's fully managed database-as-a-service. The EDB PG 15 release introduces several noteworthy features. The EDB Advanced Storage Pack improves speed of access to clustered data and heightens performance and scalability for foreign key relationships. EDB Postgres Tuner automates the database tuning process to enhance performance, drawing upon EnterpriseDB's more than 15 years of Postgres tuning expertise. EDB LDAP Sync streamlines LDAP support by eliminating the necessity to manage users in both the database and LDAP.

## SCYLLADB

This year's Database Technology Value Matrix identifies ScyllaDB as a Core Provider, acknowledging its ScyllaDB Cloud DBaaS solution, accessible via both AWS and Google Cloud. ScyllaDB serves a broad array of sectors including eCommerce, Finance,

Cybersecurity, Media, and Industrial Automation. The shard-per-core architecture of ScyllaDB is often the key factor for customers' choice; this design guarantees even distribution of storage, CPU, and RAM to each segment of the database. Features differentiate ScyllaDB include its compatibility with Cassandra and DynamoDB, pricing structure, engineering support and monitoring, and automatic backup systems. Moreover, ScyllaDB provides virtual private cloud peering and replication across multiple availability zones and has AWS reallocation features. A notable feature is its automated setup, which eliminates the need for manual configuration, subsequently reducing deployment expenses and accelerating the time-to-value.

Recent product updates and enhancements include:

- Scylla announced the release of ScyllaDB Open Source 5.2, aimed at enhancing users' experience with more robust capabilities, improved stability, and performance upgrades for the NoSQL database server and APIs. This release resolves over 100 issues, fostering a more reliable and efficient database ecosystem.
- The vendor introduced Raft-based strongly consistent schema management in ScyllaDB 5.2, an enhancement aimed at offering users a more reliable schema update process and reducing errors during concurrent schema changes. This update is projected to significantly boost system stability and overall user confidence in operations.
- Scylla announced the implementation of Raft for serialization of all schema management operations, enabling faster cluster assembly, and improved management of concurrent schema changes and updating of node IP addresses. Despite some potential limitations during node partitioning, the ability to proceed with most operations means minimized downtime and maintained data availability.
- Scylla introduced the default use of Raft for all new clusters from ScyllaDB 5.2 onwards, ensuring that every cluster-level operation requires a quorum to be executed. This strategic enhancement is designed to provide users with increased system consistency.
- Scylla announced the full adoption of Time To Live (TTL) functionality in its DynamoDB API (Alternator), granting users more control over data lifecycle management. The option to set item expiration times and configure expiration delay empowers users to manage their data storage and retrieval more effectively, leading to resource and cost optimization.
- The vendor introduced a large collection detection feature in ScyllaDB 5.2, designed to help users identify potential performance bottlenecks in advance. By allowing users to set a configurable warning threshold for "large" collections, they can better optimize their databases for maximum performance.

- Scylla announced the automation of tombstone garbage collection, replacing the need for the `gc_grace_seconds` parameter. This enhancement is set to streamline database cleanup operations for users, ensuring smoother operations and reducing the risk of unintended data resurrection.
- The vendor introduced a feature aimed at preventing timeouts when processing long tombstone sequences. This improvement is projected to reduce interruptions during data processing tasks, thus enhancing user experience and facilitating smoother execution of analytics workloads.
- Scylla announced that secondary indexes can now index collection columns, providing users with more sophisticated data retrieval options. This feature enhances the versatility of the database, empowering users with deeper insights and more complex querying capabilities.
- Scylla announced additional improvements across numerous areas including the CQL API, Amazon DynamoDB Compatible API (Alternator), Correctness, Performance and stability, Operations, Deployment and installations, Tools, Configuration, Monitoring, and tracing. These comprehensive enhancements are projected to benefit users with a more reliable database environment.

## TIGERGRAPH

TigerGraph, categorized as a Facilitator in this year's Database Technology Value Matrix, is a graph database management system that can handle large-scale data with billions of edges. Available via the cloud, TigerGraph's applications span numerous data-intensive sectors such as customer data analysis, fraud detection, IoT, and machine learning. The system utilizes C++ software and parallel processing for complex algorithm and query processing. It also features a graph query language similar to SQL and a software development kit for creating graphs and visualizations. TigerGraph primarily supports data analytics and machine learning applications.

Recent product updates and enhancements include:

- TigerGraph has announced the release of TigerGraph Cloud 3.9, its native parallel graph DBaaS, equipped with improved security and machine learning capabilities. This new version responds to the increasing demands of its growing customer base, aiming to streamline the adoption, deployment, and management of its scalable graph database platform. The engine that underpins TigerGraph Cloud is also available for on-premises or self-managed cloud installations. The platform is available as a self-managed enterprise or on fully-managed cloud services including Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure. Key upgrades in this new release include simplified streaming data ingestion, multi-



edge support, enhanced graph data science package, expanded Kubernetes functionality, and improved DevOps support. It also supports Parquet, the popular open-source big data storage format, as a data source. Further, the platform offers over 20 starter kits for industry applications such as fraud detection, supply chain analysis, and cybersecurity. Each kit is equipped with sample graph data schema, dataset, and queries focused on specific use cases.

- TigerGraph has announced its decision to support openCypher, a prevalent query language used for constructing graph database applications. In an effort to facilitate developers' transition, a limited preview translation tool has been introduced, providing insights into how openCypher support will manifest in TigerGraph's proprietary graph query language, GSQL. The move to support openCypher underlines TigerGraph's mission to enhance accessibility of graph technology. Developers familiar with openCypher can now leverage TigerGraph's performance and scalability for in-database computation and analysis of relationships within their data. TigerGraph's GSQL language offers several advantages, such as advanced querying capabilities, including 55+ algorithms available in TigerGraph's Data Science Library coded in GSQL, and enhanced performance due to easy optimization and built-in parallel processing of the underlying engine. The openCypher-to-GSQL translation tool is currently accessible to the openCypher developer community, allowing a side-by-side comparison of openCypher queries and their GSQL equivalents. In line with supporting openCypher, TigerGraph also extends its support to the industry-standard GQL. TigerGraph is part of the ISO steering committee responsible for developing GQL, the new international standard query language expected to be available in early 2024.