

WAVE REPORT

The Forrester Wave™: Vector Databases, Q3 2024

The 14 Providers That Matter Most And How They Stack Up

September 23, 2024 ● 17 min read



Noel Yuhanna
VP, Principal Analyst



With contributors:
Aaron Katz, **Sara Sjoblom**, and **Jen Barton**

Summary

In our 25-criterion evaluation of vector database providers, we identified the most significant ones and researched, analyzed, and

scored them. This report shows how each provider measures up and helps vector database professionals select the right one for their needs.

Broad Vector Capabilities, Data Management, And Performance Matter The Most

Traditional databases are unable to meet the growing demands of generative AI (genAI) due to limitations in supporting modern vector multidimensional data and performing similarity searches. Vector databases overcome these challenges by providing cutting-edge data management capabilities for storing, indexing, processing, and searching vectors efficiently. They are specifically engineered to store vector embeddings, which are numerical representations of complex data such as text, audio, images, and video. These embeddings are generated by providers like OpenAI, Hugging Face, and Cohere, with the number of dimensions varying based on the data and ML model used for vector generation. Additionally, vector databases use a variety of advanced indexing and hashing techniques, including K-dimensional trees, hierarchical navigable small world (HNSW), locality-sensitive hashing (LSH), and graph-based indexes.

Vector databases are gaining attention due to genAI, leading to a plethora of new use cases. They are crucial for providing reliable and enriched data to support genAI applications, enabling more insightful and contextually relevant responses from genAI models. Vector databases also offer an effective solution for detecting data anomalies by analyzing deviations from expected norms and providing insightful recommendations. Forrester anticipates broad growth in vector database initiatives across various industries, including financial services, retail, healthcare, manufacturing, and energy.

As a result of these trends, vector database customers should look for providers that:

1. **Support a breadth of core vector capabilities.** Unless you are building custom applications that use a vector database solely for storing and accessing vectors, you'll need comprehensive vector capabilities such as vector streaming, vector indexing, metadata management, vector search, and hybrid search. These core capabilities accelerate the development and deployment of genAI applications by automating vector operations. Look for vector database vendors offering capabilities such as advanced indexing, diverse distance metrics, real-time vector updates and deletion, broad dimensionality support, flexible metadata storage, and hybrid search.
2. **Streamline data management for vectors.** Vector databases provide a range of important data management functions for vectors, including vector storage, real-time data updates, granular access control, data integration, data semantics, resource optimization, data integrity and consistency, concurrency control, and

elastic scalability. These capabilities are critical for ensuring consistent and reliable support for vectors. When selecting a provider, look for broad data management capabilities that meet your requirements, especially when working with both vector and nonvector data. Additionally, consider vendors offering a user-friendly, persona-based UI to simplify deployments and comprehensive APIs for agile development.

3. **Deliver performance and scale at the speed of business.** Query speed is critical for customer-facing applications that leverage vector databases, especially large and complex ones that store and process tens to hundreds of millions of vectors. A vector database must handle vast volumes of vector data and scale up or down based on workload requirements. Look for vendors offering a variety of indexing algorithms, such as HNSW, IVF, and DiskANN, to meet your specific workload needs. Additionally, consider vendors that provide GPU integration, such as with NVIDIA, which can offer significant performance gains over regular CPUs. Given that most vector databases are still maturing, speak to customer references or benchmark internally when processing over 100 million vectors to ensure the solution meets your requirements.

Evaluation Summary


The Forrester Wave™ evaluation highlights Leaders, Strong Performers, Contenders, and Challengers. It's an assessment of the top vendors in the market; it doesn't represent the entire vendor landscape. You'll find more information about this market in [The Vector Databases Landscape, Q2 2024](#).

We intend this evaluation to be a starting point only and encourage clients to view product evaluations and adapt criteria weightings using the Excel-based vendor comparison tool (see Figures 1 and 2). Click the link at the beginning of this report on Forrester.com to download the tool.

THE FORRESTER WAVE™

The Forrester Wave™: Vector Databases

Q3 2024



Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.
Figure 1 - Forrester Wave™: Vector Databases, Q3 2024

		Forrester's weighting	Amazon Web Services	Couchbase	DataStax	Elastic	KX	Microsoft	MongoDB
Current offering		3.50	2.20	3.40	2.70	2.40	2.90	2.70	
Vector dimensionality	5%	5.00	1.00	3.00	1.00	3.00	1.00	1.00	
Vector embeddings	10%	3.00	3.00	5.00	3.00	3.00	1.00	3.00	
Vector metadata	5%	3.00	1.00	5.00	5.00	3.00	3.00	1.00	
Vector indexes	10%	3.00	3.00	3.00	1.00	3.00	3.00	1.00	
Vector search	5%	5.00	1.00	1.00	1.00	3.00	1.00	1.00	
Hybrid search	5%	3.00	3.00	3.00	3.00	3.00	5.00	3.00	
Stable similarity search	5%	3.00	3.00	3.00	3.00	1.00	3.00	3.00	
Search relevance	5%	3.00	1.00	5.00	3.00	3.00	3.00	3.00	
Streaming/loading vector data	5%	3.00	1.00	3.00	3.00	3.00	3.00	3.00	
Data security	5%	5.00	1.00	3.00	3.00	1.00	5.00	3.00	
Deployment	5%	1.00	3.00	5.00	3.00	1.00	1.00	3.00	
Performance	10%	3.00	1.00	3.00	3.00	5.00	3.00	3.00	
Scale-out optimization	10%	3.00	3.00	3.00	3.00	1.00	3.00	3.00	
API support	5%	5.00	3.00	3.00	3.00	1.00	3.00	5.00	
Data resilience	5%	5.00	3.00	3.00	3.00	1.00	5.00	3.00	
Database administration	5%	5.00	3.00	3.00	3.00	1.00	5.00	5.00	
Strategy		3.20	2.30	3.70	3.10	1.60	3.80	2.50	
Vision	30%	3.00	3.00	5.00	3.00	1.00	5.00	3.00	
Innovation	25%	3.00	1.00	3.00	5.00	1.00	3.00	3.00	
Roadmap	25%	3.00	3.00	3.00	1.00	3.00	3.00	1.00	
Partner ecosystem	5%	5.00	3.00	3.00	3.00	3.00	5.00	5.00	
Adoption	5%	5.00	1.00	3.00	5.00	1.00	5.00	3.00	
Pricing flexibility and transparency	5%	3.00	3.00	5.00	3.00	1.00	3.00	3.00	
Community	5%	3.00	1.00	3.00	3.00	1.00	3.00	1.00	

Partner ecosystem	5%	3.00	5.00	1.00	1.00	3.00	3.00	1.00
Adoption	5%	1.00	1.00	3.00	3.00	3.00	1.00	3.00
Pricing flexibility and transparency	5%	3.00	3.00	3.00	3.00	3.00	3.00	3.00
Community	5%	3.00	1.00	1.00	5.00	1.00	1.00	5.00
Market presence		1.00	2.20	3.80	1.80	1.80	1.00	3.80
Revenue	60%	1.00	3.00	3.00	1.00	1.00	1.00	3.00
Number of customers	40%	1.00	1.00	5.00	3.00	3.00	1.00	5.00

All scores are based on a scale of 1 (weak) to 5 (strong).

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.
 Figure 2 - Forrester Wave™: Vector Databases Scorecard, Q3 2024

Vendor Offerings

Forrester evaluated the offerings listed below (see Figure 3).

Vendor	Product evaluated
Amazon Web Services	Amazon Aurora for PostgreSQL, Amazon Relational Database Service [RDS] for PostgreSQL, Amazon OpenSearch Service, Amazon MemoryDB, Amazon DocumentDB, Amazon Neptune Analytics, Amazon Knowledge Bases for Amazon Bedrock, AWS GuardDuty
Couchbase	Vector Search
DataStax	Astra DB Vector
Elastic	Elastic Search
KX	KDB.AI
Microsoft	Azure AI Search, Azure Cosmos DB, Azure Database for PostgreSQL — Flexible Server
MongoDB	MongoDB Atlas Vector Search
Neo4j	Neo4j Graph Database
Oracle	AI Vector Search in Oracle Database
Pinecone	Pinecone serverless
Qdrant	Qdrant
Redis	Red Vector Database
SingleStore	SingleStore
Zilliz	Zilliz Cloud/Milvus

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

Figure 3 - Evaluated Vendors And Product Information

Vendor Profiles

Our analysis uncovered the following strengths and weaknesses of individual vendors.

Leaders

1. **Zilliz delivers a cutting-edge, high-performance database designed for cloud scalability.** Known as the creator of Milvus, the leading open-source vector database, Zilliz excels in managing vast amounts of vector data. It leverages optimized storage, efficient management, and search capabilities. Zilliz is at the forefront of innovation, delivering exceptional speed and efficiency in vector processing and search to support real-time AI applications. With a customer-centric approach, Zilliz integrates feedback from its robust community to continuously enhance its roadmap with new features and functionality. Zilliz's superior roadmap focuses on distributed index, advanced reranking of results, administration, broader data security certifications, multicloud support, and data intelligence.

Zilliz is strong in several areas, including vector dimensionality, vector indexing, performance, and scalability. However, it lags in administration tooling and data security certifications. Reference customers praise Zilliz's strong open-source foundation, its reasonable cost, and its speed, technical support, reliability, and scalability. Zilliz is particularly well suited for customers prioritizing high-end performance and low-latency access to extensive vector embedding data for advanced AI applications.

View [Zilliz's detailed scorecard](#).

2. **DataStax has a strong vision with broad vector capabilities.** DataStax has been instrumental in developing and commercializing Apache Cassandra, a highly scalable, distributed database. Building on this foundation, DataStax AstraDB provides a viable vector database to support modern AI applications. The vendor is actively expanding its vector database capabilities, particularly focusing on genAI, RagStack, serverless, and distributed architecture. DataStax's superior vision emphasizes a simplified and automated genAI stack, enhanced data ingestion, advanced developer tooling, and search relevance to empower developers in creating modern retrieval-augmented generation (RAG) and genAI applications. DataStax's roadmap emphasizes strengthening capabilities in enhanced security capabilities and certifications, improvements in performance and scale, GPU support, knowledge graph, and advanced automation.

DataStax excels in vector embeddings, vector metadata, search relevance, and deployment options. It needs to improve on vector search. Reference customers appreciate DataStax for its strong technical support, reliability, flexibility, search relevance, consistency, and relatively low total cost of ownership. Some customers express concerns about support for additional vector embeddings and performance challenges, and say they want expanded hybrid search capabilities. DataStax is particularly well suited for customers leveraging Apache Cassandra, where it offers seamlessly integrated vector search capabilities, or those seeking a comprehensive data platform for developing genAI

applications, including vector search, integration, and data management.

View [DataStax's detailed scorecard](#).

Strong Performers

1. **AWS has several viable vector databases but they are limited to AWS cloud.**

Known for its extensive range of cloud databases trusted by tens of thousands of organizations for mission-critical applications and insights, Amazon Web Services (AWS) has extended its offerings to include vector support. This expansion introduces new capabilities such as vector search, enhanced vector security, and efficient vector processing and administration. AWS boasts a competitive roadmap and vector database strategy, and its vision and innovation are on par with others evaluated. AWS's roadmap is centered around supporting genAI workloads, integration with AWS services, and delivery of advanced automation.

AWS has strong capabilities in vector dimensionality, database administration, API support, data security, and vector search. AWS lags in multiple cloud deployment options. Reference customers value AWS's broad database ecosystem, support for vectors, expanded API capabilities, and data versatility. Only one of AWS's reference customers responded to Forrester's outreach for this evaluation. AWS is an ideal choice for those looking to utilize existing AWS databases to efficiently manage vector data without extra cost or a steep learning curve.

View [AWS's detailed scorecard](#).

- 2. Microsoft has a promising vector solution but lacks advanced vector capabilities.** Microsoft offers multiple databases, including Azure CosmosDB, Azure Database for PostgreSQL, and Azure SQL database, each with vector capabilities tailored to support specific use cases and workloads. These databases seamlessly integrate with Azure OpenAI and other AI services for model deployment and AI application development. Microsoft's strong vision prioritizes automating, accelerating, and democratizing data through data fabric and AI, harnessing the power of vector search. Microsoft's roadmap and innovation are on par with other competitors. Its roadmap emphasizes automation, integration with Azure fabric, enhanced embeddings, improved performance, and refined reranking strategies.

Microsoft stands out in hybrid search, data security, data resilience, and robust database administration. Where Microsoft lags is in vector dimensionality, vector embeddings, and vector search capabilities. Reference customers like Microsoft's strength in hybrid search, its low latency, its support for vector search, its consistency, and its suitability for genAI workloads. Some have concerns regarding costs, API versioning, relevance scoring, geo-replication, and tooling. Microsoft is particularly suitable for organizations seeking to leverage vector capabilities integrated with Azure's data platform to meet the demands of modern genAI applications.

View [Microsoft's detailed scorecard](#).

3. **Oracle embraces the vector database trend but its solution is not yet mature.**

Despite being a latecomer in the vector database space, Oracle has aggressively ramped up its offering, challenging established vendors. Oracle integrates vector capabilities directly into its core database engine, providing a streamlined experience for developing modern genAI applications. Oracle boasts a competitive roadmap and vector database strategy, and its vision and innovation are on par with others evaluated. Oracle's roadmap emphasizes expanding capabilities in advanced vector search, native hardware acceleration, broader LLM integration, advanced natural language processing, and highly scalable indexing.

Oracle is strong in vector metadata, data resilience, streaming/loading vector data, data security, deployment options, and database administration. All other functionalities are on par with others. Reference customers like Oracle's solution for its operational ease, friendliness for genAI use cases, technical support, and SQL integration. They also like that it's part of Oracle's ecosystem, that it's built into the core engine, and the ready access to product managers. Oracle is an excellent choice for customers seeking a multimodel database with vector capabilities to support modern genAI use cases or for existing Oracle customers.

View [Oracle's detailed scorecard](#).

4. **Neo4j has lots of potential with GraphRAG but lacks customer usage.** Neo4j is a leading graph database provider specializing in knowledge graph use cases. It excels at simplifying complex data relationships, pattern recognition, and

delivering insights from connected data. Recently, Neo4j introduced vector search, blending graph database capabilities with vector search to enhance AI and semantic search applications. Neo4j's strong vision is clear and differentiated, focusing on delivering a multimodel, multicloud, and unified platform for genAI applications by leveraging GraphRAG to provide context and accuracy. Neo4j's roadmap prioritizes enhancing capabilities in knowledge graphs for GraphRAG, automation, optimizing vector storage, improving performance and scalability, and expanding the partner ecosystem and tooling.

Neo4j's GraphRAG approach leads to standout search relevance. However, it lags in vector dimensionality, vector indexes, and scale-out. Reference customers like Neo4j's flexibility to support structured and unstructured data, its graph data science modules, and its technical support. However, some have concerns about lack of clear transparency on how vector metrics are calculated and that there's no native integration with other databases. Neo4j only provided one reference customer for this evaluation. Neo4j is a good fit for customers looking to blend knowledge graphs with vector search capabilities to support RAG applications, advanced AI applications, and enhanced recommendation engines.

View [Neo4j's detailed scorecard](#).

5. Elastic has a customizable offering but needs advanced vector capabilities.

Elastic is a popular search engine built on Lucene, offering a distributed, multitenant, and full-text RESTful search capability. For a few years, it has supported vectors, enabling organizations to support modern genAI

applications. Elastic is also used for observability, security, and custom application building. Elastic's early and ongoing investments form the foundation of its vision of vector search and database. Elastic's vision aims for a seamlessly integrated search experience designed to accelerate diverse genAI use cases for developers. Elastic can improve its roadmap with advanced capabilities around automation and security certifications, and support more vector indexes.

Elastic's standout capability is the storing, management, and integration of vector metadata. Elastic needs to improve around more advanced vector capabilities like vector dimensionality, vector indexes, and vector search. Reference customers like Elastic's technical support and open-source support. However, some customers have called out concerns around cost, horizontal scale, uptime/availability, better conversion of PDFs, speed of indexing, and limited support for ingesting content. Elastic is a good fit for customers who are looking to leverage heavily customized vector search within modern AI applications.

View [Elastic's detailed scorecard](#).

- 6. Pinecone has enterprise scale but does not have comprehensive vector capabilities.** Pinecone is a native vector database founded in 2019, dedicated to helping organizations build and support modern AI applications. Initially, the company aimed to make the database more accessible to engineering and development professionals, simplifying the management of vector data. Over the years, Pinecone has delivered cutting-edge innovation and maintained a

strong commitment to vector databases. Pinecone's vision is to simplify AI in production, prioritizing a developer-first experience. Its compelling roadmap focuses on API versioning, enhanced data security, improved storage integration, serverless automation, and performance.

Pinecone excels in scale-out optimization and vector dimensionality. Despite being a native vector database, it lags this competitive group in vector metadata, vector indexes, vector search, hybrid search, data security, and API support. Reference customers like Pinecone's support, competitive pricing, high-scale capabilities, and comprehensive documentation. Nevertheless, some customers have raised concerns about availability, reliability, and unsatisfactory service-level agreement (SLA) guarantees. Pinecone is a good fit for customers seeking a native vector database that can manage large volumes of vector embeddings within a scale-out architecture.

View [Pinecone's detailed scorecard](#).

7. **Qdrant has a viable vector database but lags in advanced database capabilities.** Founded in 2021, Qdrant is a native open-source vector database designed to offer flexibility and ease of use, helping developers build applications that rely on vector similarity search. Qdrant's vision, innovation, and roadmap are on par with those evaluated. The company is poised to enhance its vector capabilities, bolstering its competitive position. The vendor's roadmap includes expanding capabilities on advanced hybrid search, automation, and performance improvements.

Qdrant excels in supporting high vector dimensionality, enabling embeddings for very complex data. It needs to improve on vector indexes, data security, and API support. Reference customers like Qdrant's geo-spatial search, robustness, API integration, documentation, support and services, and simplicity. However, some have concerns about scalability, monitoring, hybrid search complexity, and disaster recovery. Qdrant is a good fit for customers seeking a native vector database with broad vector dimensionality capabilities for highly complex data that is easy to deploy for supporting modern genAI/LM apps, semantic and similarity searches, and data science.

View [Qdrant's detailed scorecard](#).

Contenders

1. **MongoDB extends its database to support vectors but must increase innovation.** MongoDB is a widely used document database featuring an integrated suite of services that enable development teams to build and support modern business applications. In June 2023, MongoDB introduced vector capabilities with MongoDB Atlas vector search, allowing organizations to leverage vectors with a minimal learning curve. MongoDB's vision and innovation are on par with those evaluated. MongoDB can improve its roadmap by offering more advanced vector capabilities and automation.

MongoDB's strength lies in API support, enabling developers to interact via a variety of popular languages, as well as database administration. However, it

lags in vector dimensionality, vector metadata, vector indexes, and vector search. Reference customers like MongoDB's expansion of the product to support vector embeddings without requiring developers to be deeply knowledgeable about vectors. However, there are concerns about comprehensive documentation, customer support, and production bugs. Only one of MongoDB's reference customers responded to Forrester's outreach for this evaluation. MongoDB is well suited for customers who are already using a document database and want to integrate vector capabilities. It is also ideal for those seeking a unified multimodel database that supports both structured and unstructured data for RAG and AI applications.

View [MongoDB's detailed scorecard](#).

- 2. Redis offers a viable vector database but lags on advanced vector capabilities.** Redis is a multimodel, open-source, and in-memory database that supports documents, graphs, time services, and object data models for building modern applications and insights. In August 2023, with the release of Redis 7.2, the platform introduced vector search capabilities, enabling the storage, indexing, and searching of vector embeddings. While Redis' vision and roadmap are on par with other vendors, it has yet to demonstrate significant innovation in the vector space. Moving forward, Redis' roadmap focuses on automation, expanding vector search functionality, and improving ease of use.

Redis is on par with others evaluated in some criteria, including vector dimensionality, vector metadata, vector search, streaming/loading vector data, data security, deployment, data resilience, scale-out optimization, and database

administration. Several features, including search relevance and performance, lag this evaluative set. Reference customers like Redis' responsiveness, knowledgeable solutions engineers, support for real-time data, ease of use, and documentation. However, some concerns exist around cost, partner ecosystem, and tooling. Redis is a good fit for customers who want a multimodel database that can support vectors in a cached layer for faster processing.

View [Redis' detailed scorecard](#).

3. **SingleStore has a viable, but immature, vector offering.** SingleStore is a versatile translytical database designed to seamlessly handle both transactional and analytical workloads within a unified platform. SingleStore excels in real-time ingestion, processing, and analytics, leveraging cloud-native features and offering SQL-compatibility. Notably, SingleStore introduced vector processing with its recent version 8.5 release. SingleStore lags in innovation and could improve it by partnering with application and infrastructure providers that offer tools integrated with the database. Additionally, the vendor needs to expand its vector capabilities to target not only developers but also the broader business community and foster broader data and AI partnerships. Looking ahead, the vendor's roadmap includes expanding capabilities for advanced automation, GPU integration, comprehensive search functionalities, broader data security capabilities, and improved performance.

SingleStore provides excellent data resilience capabilities for use cases where data loss or disruption are unacceptable. However, it lags in core vector capabilities like vector search as well as on data security. Reference customers

like SingleStore's ability to seamlessly process vectors alongside traditional data, its user-friendly interface, dedicated support services, and its integration of vector, column store, and relational data. However, references have some concerns about cost, the completeness of documentation, product maturity, and the pace of innovation. SingleStore is particularly well suited for customers seeking a unified platform for transactional, operational, and analytical workloads along with vector search capabilities.

View [SingleStore's detailed scorecard](#).

4. **Couchbase offers an easy-to-use vector but lacks comprehensive capabilities.**

Couchbase is known for its multimodel database, adeptly combining key-value and document data to cater to a wide range of application needs. Recently, Couchbase introduced support for vector search, empowering developers to manage genAI and RAG workloads efficiently. Couchbase's vision and roadmap are on par with others. Its vision could be further enhanced by incorporating advanced vector capabilities, which would accelerate use cases and broaden the platform's accessibility. Couchbase's roadmap includes expanding capabilities on automation, ease of use, data security, expanded support for vector embedding and search capabilities, and optimizations for performance and scalability.

Couchbase did not demonstrate any category leading capabilities, mostly rating on par with the competition. However, Couchbase lags in delivering comprehensive vector capabilities like vector dimensionality, vector metadata, vector search, performance, and data security. Couchbase did not provide

reference customers for this evaluation. Couchbase is a good fit for customers looking for a multimodel database with vector capabilities to support both structured and unstructured data for genAI and RAG applications.

View [Couchbase's detailed scorecard](#).

5. **KX offers a high-performance database but needs to improve its vector capabilities.** KX is known for its high-performance, low-latency data, and analytics platforms. At the core of its offerings is kdb+, an in-memory, time series database optimized for managing large datasets. Recently, KX introduced KDB.AI, which integrates vector storage, processing, and search capabilities to support modern AI applications such as vector and similarity searches. Its kdb+ uses parallel processing with GPUs through NVIDIA CUDA integration and other GPU libraries. KX's strategy is to focus on high performance, but the company needs to expand its vision to compete with vendors offering broader vector data capabilities. The vendor's roadmap focuses on advancing search, expanding its ecosystem, enhancing developer productivity, boosting performance metrics, and implementing automation to support broader use cases.

KX delivers strong tuning capabilities to deliver optimal performance. However, it lacks data security, scale-out, API support, and data resilience. Reference customers like KX's performance, reliable technical support, interoperability, seamless integration capabilities, and product reliability. However, there are some concerns regarding adoption rates, community size, long-term sustainability, third-party integration, and ecosystem development. KX is a good fit for customers who want high-performance, low-latency access to vector data,

leveraging optimized hardware architectures including GPUs, and distributed memory for critical AI and RAG applications.

View [KX's detailed scorecard](#).

Evaluation Overview

We grouped our evaluation criteria into three high-level categories:

1. **Current offering.** Each vendor's position on the vertical axis of the Forrester Wave graphic indicates the strength of its current offering.
2. **Strategy.** Placement on the horizontal axis indicates the strength of the vendors' strategies, including elements such as vision and innovation.
3. **Market presence.** The size of each vendor's marker on the graphic reflects Forrester's assessment of its market presence.

Vendor Inclusion Criteria

Each of the vendors we included in this assessment has:

1. **Enterprise vector database capabilities.** Invited vendors have a vector database solution that includes features such as vector search, vector-based indexes, vector data security, vector data compression, LLM integration, real-

time vector query, scale-out optimization, hybrid search, vector data processing, and genAI optimization.

2. **Support for a broad set of use cases.** The vector database solution has customers using the solution in a production environment for use cases such as RAG, recommendation engines, genAI/LLM apps, semantic and similarity search, and CX personalization.
3. **Forrester mindshare.** Included vendors have been significantly mentioned through Forrester enterprise client interactions calls during the past 12 months related to vector database topics.

Supplemental Material

Online Resource

We publish all our Forrester Wave scores and weightings in an Excel file that provides detailed product evaluations and customizable rankings; download this tool by clicking the link at the beginning of this report on Forrester.com. We intend these scores and default weightings to serve only as a starting point and encourage readers to adapt the weightings to fit their individual needs.

The Forrester Wave Methodology

A Forrester Wave is a guide for buyers considering their purchasing options in a technology marketplace. To offer an equitable process for all participants, Forrester follows The Forrester Wave™ Methodology to evaluate participating vendors.

In our review, we conduct primary research to develop a list of vendors to consider for the evaluation. From that initial pool of vendors, we narrow our final list based on the inclusion criteria. We then gather details of product and strategy through a detailed questionnaire, demos/briefings, and customer reference surveys/interviews. We use those inputs, along with the analyst's experience and expertise in the marketplace, to score vendors, using a relative rating system that compares each vendor against the others in the evaluation.

We include the Forrester Wave publishing date (quarter and year) clearly in the title of each Forrester Wave report. We evaluated the vendors participating in this Forrester Wave using materials they provided to us by June 11, 2024 and did not allow additional information after that point. We encourage readers to evaluate how the market and vendor offerings change over time.

In accordance with our vendor review policy, Forrester asks vendors to review our findings prior to publishing to check for accuracy. Vendors marked as nonparticipating vendors in the Forrester Wave graphic met our defined inclusion criteria but declined to participate in or contributed only partially to the evaluation. We score these vendors in accordance with our vendor participation policy and publish their positioning along with those of the participating vendors.

Integrity Policy

We conduct all our research, including Forrester Wave evaluations, in accordance with the [integrity policy](#) posted on our website.

About Forrester Reprints <https://go.forrester.com/research/reprints/>

© 2024. Forrester Research, Inc. and/or its subsidiaries. All rights reserved.