Everest Group®

Everest Group Innovation Watch: Vector Data Platforms

Focus on VAST Data

July 2026

Copyright © 2026 Everest Global, Inc. | this document has been licensed to **VAST Data**

Research background

As enterprises accelerate AI adoption, managing and retrieving high-dimensional data has become a critical capability. Traditional databases are increasingly complemented and augmented by vector data platforms designed specifically to store and query embeddings generated by AI and machine learning models. These platforms support advanced applications, such as semantic search, recommendation systems, and Retrieval-augmented Generation (RAG), enabling efficient handling of unstructured and multimodal data.

With rising data complexity and real-time retrieval demands, vector data platforms are gaining prominence for their ability to combine vector storage, search, analytics, and hybrid querying within a unified architecture. Their scalable, low-latency design and integration with AI pipelines make them foundational to next-generation intelligent applications.

In this study, we assess leading providers of vector data platforms that enable embedding-based data management and retrieval. The assessment evaluates each provider's market performance, scalability, integration flexibility, and innovation maturity, using the Everest Group Innovation Watch Assessment framework. Providers are categorized into four segments: Luminaries, Fast Followers, Influencers, and Seekers based on their innovation and market potential.

The report includes:

- A comparative assessment of leading vector data platform providers
- An overview of each provider's platform capabilities, architecture, and product maturity
- Insights into how vector data platforms are addressing key enterprise challenges in embedding-based data management and search

Scope of this report

Geography: global

Providers: Amazon, Chroma, Clickhouse, Cloudflare, Cockroach Labs, Couchbase, Databricks, Elastic, Google, IBM, InterSystems, KX, Microsoft, MongoDB, Oracle, Pinecone, Progress, Qdrant, Redis, SAP, SingleStore, Snowflake, Supabase, Teradata, VAST Data, Vespa.ai, Weaviate, and Zilliz

Product(s): vector data platforms

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Scope of the assessment

This assessment evaluates vector data platforms that enable production-grade vector storage, indexing, similarity search, and retrieval to support a wide range of workloads, such as Retrieval-Augmented Generation (RAG), semantic search, agent memory, recommendations, and anomaly detection.

As part of the offerings assessed, the vector data platform must:

- Store high-dimensional vectors with persistent, durable storage across deployment models
- Enable efficient indexing and the retrieval of vectors
- Operate as a production-ready platform or database with enterprise-grade security, governance, defined service level objectives (SLOs) and agreements (SLAs)
- Integrate with broader AI/ML ecosystems, offering connectors and tooling that support diverse workloads and use cases

Vector data platforms categories evaluated

XX Focus of the study

1



Native vector database platforms

Providers whose core product is a purpose-built vector database with specialized indexing, storage, and retrieval for high-dimensional data

Example: an independent platform created to efficiently store embeddings for millions of items with advanced similarity search

2



Traditional DBMSs with vector features

Established database providers (relational, NoSQL, and others) that have added support for native vector data types, ANN indexes, and retrieval capabilities

Example: a traditional SQL database that extends its engine to allow vector columns and hybrid queries

3



Vector-embedded cloud and data platforms

Cloud providers offering managed, scalable vector storage and retrieval as part of their broader data and AI portfolio

Example: a product part of a broader cloud platform that allows enterprises to upload embeddings and query them at scale with built-in autoscaling and global availability

5



Enterprise search platforms

Providers focused on enterprise or multimodal search platforms where the primary product is search, not a stand-alone vector database

Example: an enterprise knowledge search platform that adds vector ranking as one feature alongside keyword search and document indexing

4



Vector-search-only platforms

Libraries, SDKs, and model-centric services that provide vector search but not full database functionality

Example: an open-source toolkit that developers embed directly into applications to run nearest-neighbor searches

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Overview of the vector data platforms and the value chain



Ingestion

Enterprise data types and sources
(text, images, videos, audio, documents, PDFs, etc.)



Vector data platform

Embedding model
Transforms unstructured data into numerical vectors

Retrieval function
Measures the similarity between user query and indexed data

Indexing function
Stores the embeddings in a multidimensional space, based on similarity and relevance between the nodes (embeddings)

Vector search

Storage engine
Vector/Graph/Relational/NoSQL/Others



Enablers | API ecosystem | Cloud | LLM model

From data to intelligent applications

- Vector data platforms enable intelligent applications by transforming multimodal data (text, images, audio, video, and documents) into vector embeddings, thereby expanding the range of enterprise data that can be effectively used
- The vector data platforms then manage these through three core functions:
  - **Embedding models** turn unstructured inputs into vector representations
  - **Indexing functions** store embeddings in a multidimensional space, enabling efficient organization and fast access
  - **Retrieval functions** identify the most similar vectors to a given query, delivering relevant results
- The vectors are stored in underlying engines (relational, NoSQL, graph, or purpose-built vector stores), ensuring durability and persistence
- Broader adoption is supported by APIs, cloud infrastructure, and LLM ecosystems that integrate vector data platforms into enterprise workflows
- Key enterprise applications of vector data platforms include recommendation engines, Retrieval-Augmented Generation (RAG), and multimodal search

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

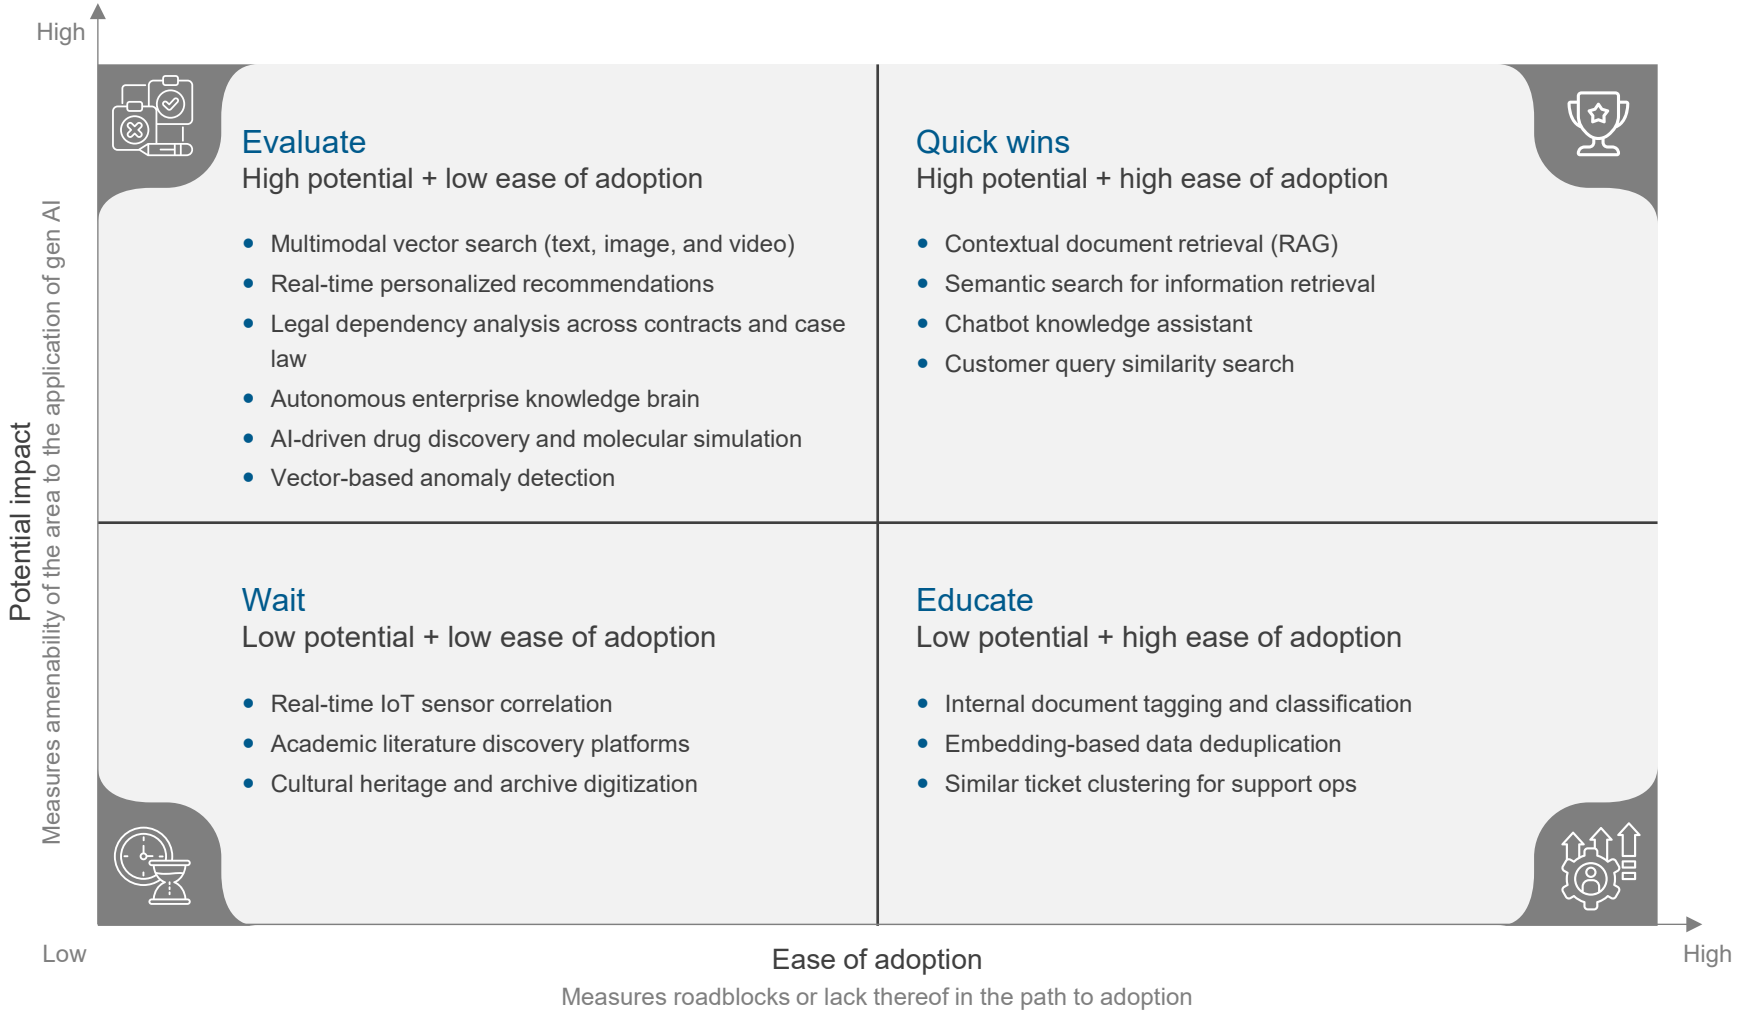
Representative challenges by business function, and outcomes enabled by vector data platforms

Function	Key business problems and challenges	Outcomes with vector data platforms
 IT and IT infrastructure	• Incorrect responses from gen AI-based search tools and chatbots • Underutilized enterprise data, need for manual tagging and metadata maintenance • Slow incident resolution due to poor discovery of similar historical incidents	• Contextually aware and more accurate responses using RAG • Scalable vector databases with large embedding stores, inclusion of multimodal data through vector embeddings, and automated content clustering and tagging • Use semantic search to identify similar historical incidents using embeddings across tickets, logs, and runbooks
 Finance and accounting	• Duplicate vendors and invoices cause leakage and fraud risk • Tracking down transaction differences and contract terms • Spend data that cannot be meaningfully categorized • Ineffective fraud detection for **unknown** fraud patterns	• Nearest-neighbor search to identify duplicates • Vector search for conditions and offered terms in past records • Vector search for semantic spend classification based on embeddings of spend line items and category definitions • Identify outliers and anomalous neighbors in vector space using transaction behaviors represented in vectors
 Customer support and contact centers	• Long response time for query resolution • Inconsistent support accuracy across agents • Difficulty handling large ticket volumes • Poor self-service experience • Free-text user complaints are hard to analyze	• Vector search for similar past cases and solutions • Real-time contextual assistance to agents • Intelligent ticket routing and summarization • Conversational self-service powered by RAG • Sentiment analysis and clustering for better and faster response

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Vector data platforms use case prioritization

Evaluating maturity, business value, and ease of integration for enterprise adoption



Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Vector data platforms' impact on solving business problems (page 1 of 2)

Implement scalable AI use cases through vectorized data foundations, hybrid retrieval, and closed loop feedback

Use case	Value addition from gen AI	Challenges in adoption	Keys to success
Chatbots enhanced with RAG	• Contextually relevant, accurate, and timely answers • Cost-effective approach for introducing domain-specific knowledge	• Quality and accuracy of output dependent heavily on knowledge base • Introduces system complexity across ingestion, retrieval, and orchestration	• Maintenance of knowledge base with the latest information needed for the purpose • Effective vector mapping and embedding with the keywords essential for purpose
Visual search in retail and e-commerce	• Faster search process using Approximate Nearest Neighbor (ANN) indexing • More efficient in producing similar results compared to keyword descriptors	• High dependency on upstream image governance and standards • Complex ETL pipeline needed to synchronize with metadata	• Standardized image normalization as part of the vector pipeline • Integrating with hybrid filtering for better user experience
Real-time personalized recommendation	• Recommendations populated even for new users with limited history • Suggestions based on content similarity and recent preferences; not just keywords, titles, and descriptors • Adjacent suggestions to increase time spent per user and understand user preferences	• Handling high velocity data streams requires distributed streaming frameworks and tools • High computational power and resources required for managing latency and instant recommendations • Balance needed for system reliability and load balancing	• Integrating customer reviews and engagement signals into recommendation parameters • Implement event-driven pipelines to continuously update user and item representations • Integrate domain-aware vector embeddings with context-enriched recommendation logic

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Vector data platforms' impact on solving business problems (page 2 of 2)

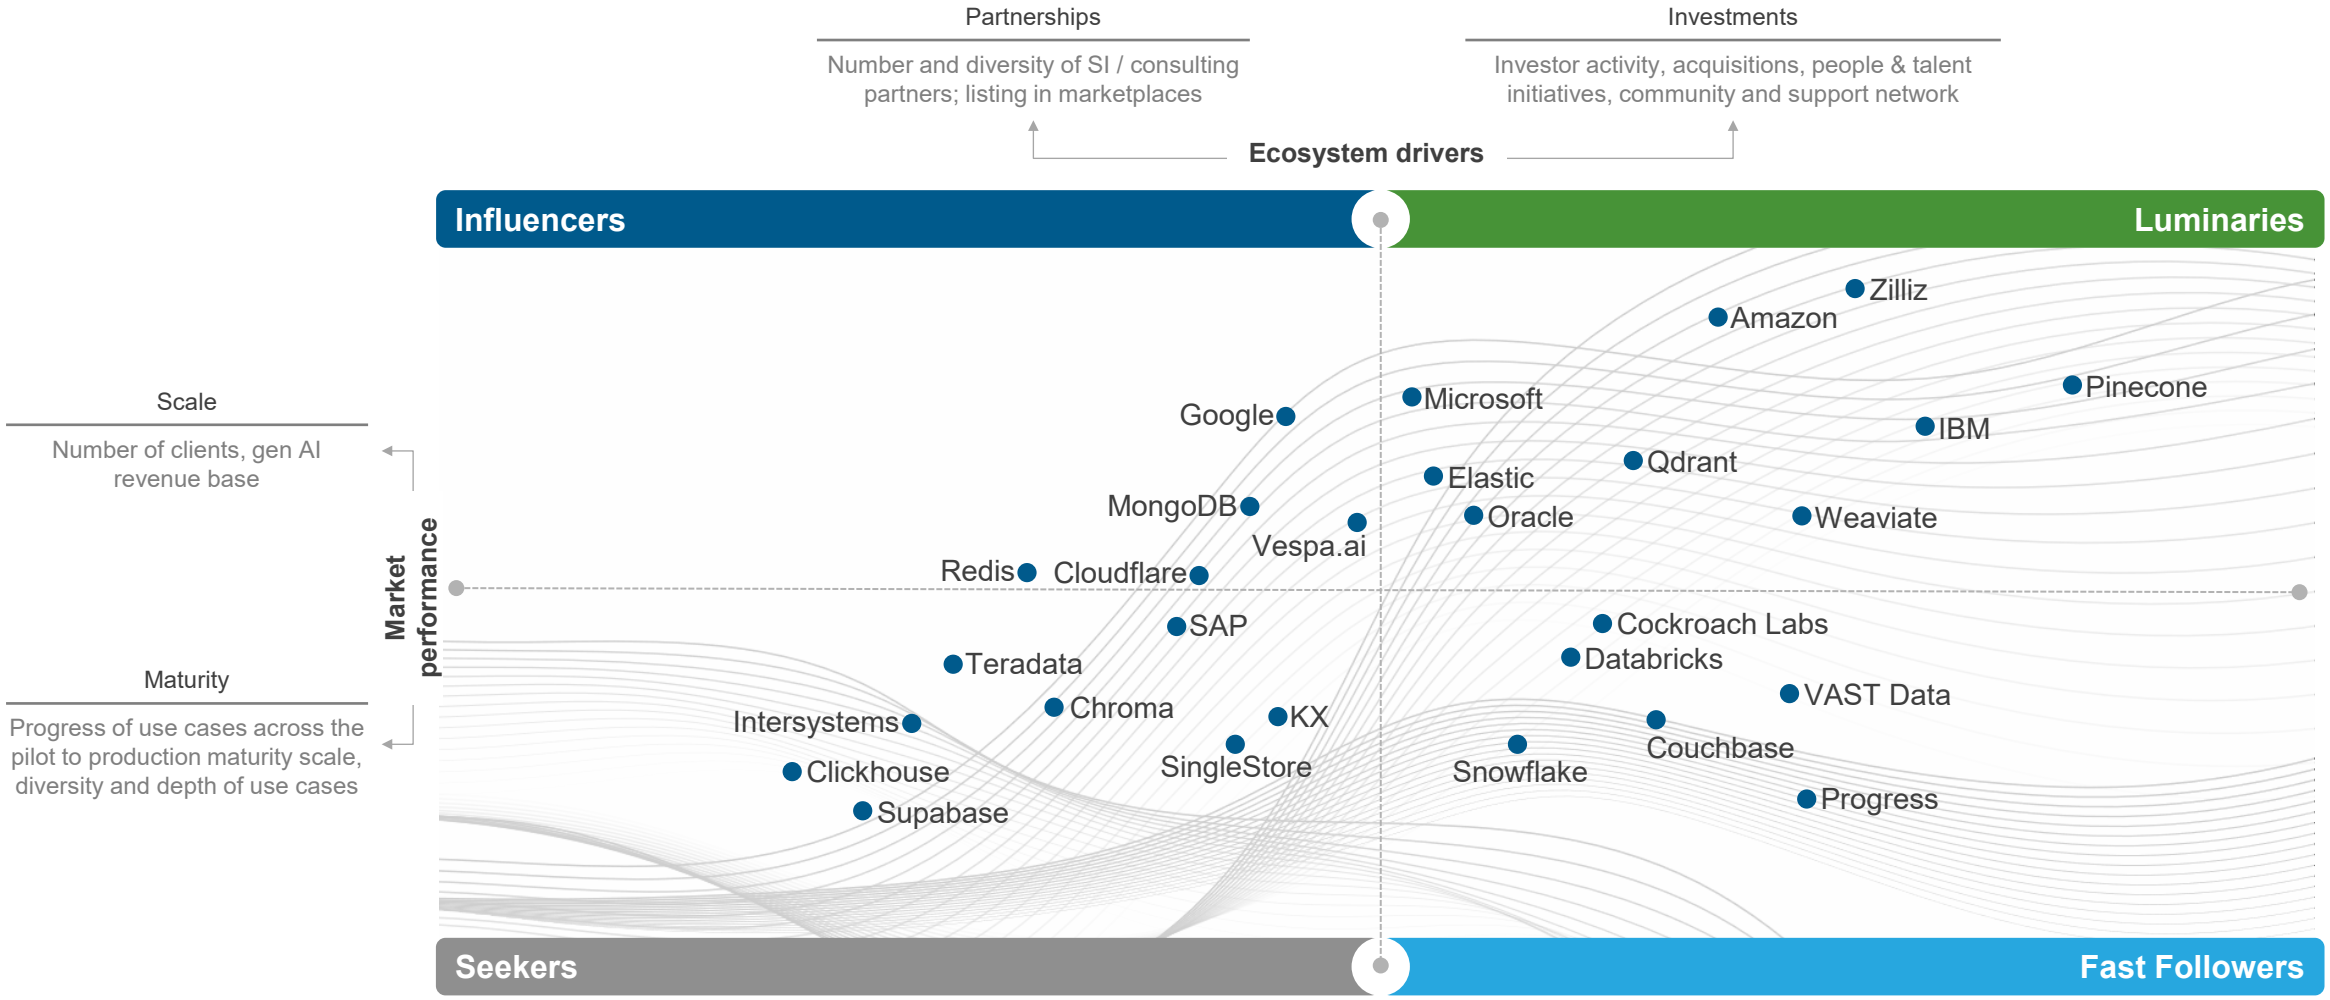
Operationalize vector capabilities through domain-aware model training, continuous monitoring, and human-in-the-loop validation

Use case	Value addition from gen AI	Challenges in adoption	Keys to success
 Anomaly and fraud detection	• Vectors encode normal behavior patterns across high-dimensional signals • Identification of novel fraud patterns that do not match known rules or labeled examples • Horizontal scalability to accommodate vast amount of growing data sets	• High-dimensionality increases computational complexity and reduced accuracy • Detecting noise as anomalies due to overfitting of data in high dimensions • Evolving fraud patterns and techniques create concept drift	• Implement human-in-the-loop workflows for investigation and feedback • Continuously update embeddings to capture evolving patterns • Continuous monitoring for detecting and addressing concept drift anomalies
 Sentiment analysis	• Embeddings capture context, tone, and intent, enabling detection of varied sentiments • Consistent sentiment analysis across languages and channels	• Context and domain-specific jargon can cause wrong results • Sarcasm-driven misclassification reduces sentiment accuracy and limits confidence in analysis	• Model training should include domain-specific jargon and context for accuracy • Flag low-confidence or ambiguous (sarcastic) sentiment and build context using human-in-the-loop validation
 Similarity search for drug discovery and retrosynthetic design	• Faster molecule and interaction detection • Real-time insights and increased automation of process • Effective synthesis planning	• High dimensionality and complexity of chemical and biological data • Integration with legacy cheminformatics tools and workflows • High compute cost for large-scale operations	• Use verified and vetted domain-specific molecular and reaction embeddings • Combine vector similarity with chemical constraints and rules and domain-specific filters • Optimize resource utilization and support dynamic scaling

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Everest Group Innovation Watch Assessment

Everest Group assessed 28 leading vector data platform providers on their market performance and ecosystem drivers, and classified them into Luminaries, Influencers, Fast Followers, and Seekers | VAST Data is positioned as a Fast Follower



 Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Summary dashboard | market performance and ecosystem drivers' assessment for vector data platforms (page 1 of 4)

Measure of capability: Low High

	Market performance			Ecosystem drivers		
Provider	Market adoption	Maturity	Overall	Partnerships	Investments	Overall
Amazon						
Chroma						
Clickhouse						
Cloudflare						
Cockroach Labs						
Couchbase						
Databricks						

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Summary dashboard | market performance and ecosystem drivers' assessment for vector data platforms (page 2 of 4)

Measure of capability: Low High

	Market performance			Ecosystem drivers		
Provider	Market adoption	Maturity	Overall	Partnerships	Investments	Overall
Elastic						
Google						
IBM						
InterSystems						
KX						
Microsoft						
MongoDB						

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Summary dashboard | market performance and ecosystem drivers' assessment for vector data platforms (page 3 of 4)

Measure of capability: Low High

	Market performance			Ecosystem drivers		
Provider	Market adoption	Maturity	Overall	Partnerships	Investments	Overall
Oracle						
Pinecone						
Progress						
Qdrant						
Redis						
SAP						
SingleStore						

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Summary dashboard | market performance and ecosystem drivers' assessment for vector data platforms (page 4 of 4)

Measure of capability: Low High

	Market performance			Ecosystem drivers		
Provider	Market adoption	Maturity	Overall	Partnerships	Investments	Overall
Snowflake						
Supabase						
Teradata						
VAST Data						
Vespa.ai						
Weaviate						
Zilliz						

Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

VAST Data | overview

Company overview

VAST Data is an AI infrastructure company that delivers the VAST AI Operating System, a unified data platform designed for data-intensive applications and agentic computing. It provides integrated storage, database, and compute capabilities through a software platform built on the Disaggregated, Shared-Everything (DASE) architecture. It supports hyperscale AI workloads, vector data, and real-time data processing.

Headquarters: New York, New York

Website: www.vastdata.com

Key clients

- Chan Zuckerberg Initiative
- CoreWeave
- Lambda
- Pixar

Key partnerships

- AMD
- Arista
- AWS
- Dell
- Google Cloud
- Kaytus
- NVIDIA
- xFusion

Recent deals/investments/announcements

- 2025: partnered with Lablup, a developer of the Backend.AI platform, to integrate the VAST AI Operating System and deliver a high-performance data foundation for Korea's Sovereign AI initiative, resolving infrastructure bottlenecks and enabling large-scale domestic model training across national research environments
- 2025: partnered with Microsoft, Google, and AWS to deploy the VAST AI OS natively on Azure, GCP, and AWS. This architecture enables hybrid and multi-cloud agentic AI workloads by providing unified data services and high-performance orchestration that integrates directly with hyperscaler GPU-accelerated infrastructure
- 2025: partnered with the Texas Advanced Computing Center (TACC), a leading academic supercomputing institution, to provide the VAST AI Operating System as the data layer for Horizon, the NSF's flagship supercomputer, extending a multi-year collaboration spanning Stampede3 and Vista to accelerate large-scale scientific and AI research

Key market segments served

Industries

- CSPs
- Energy and Utilities
- Financial Services
- Healthcare & Life Sciences
- Media & Entertainment

Commercial model

Subscription-based	User-based	Consumption-based	Outcome-based
✓		✓	

Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

VAST Data | offerings summary

Database offering name	Current version	Release year	Description
VAST AI Operating System	VAST AI OS 5.4	2025	It is built on the Disaggregated, Shared-Everything (DASE) architecture, which separates compute and storage for independent scalability. It uses CNodes for stateless computation and DNodes for persistent, high-performance storage. It enables parallel, shard-less querying for vector data. It supports real-time indexing and retrieval, making it ideal for RAG pipelines and large-scale vector workloads.

Platform specifications

Deployment options: managed and self-hosted

Cloud platforms supported: AWS/Azure/GCP

Compliance protocols supported: SOC 2, ISO 27001, HIPAA, and GDPR

Free tier/trials: VAST Data Labs for experimentation

Index types supported: proprietary hierarchical clustering

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

VAST Data | case studies

Client name/description	Business challenge	Solution	Impact
Leidos	Leidos faced overwhelming cybersecurity workloads as trillions of daily events exceeded human capacity for triage. It created alert fatigue and slowed response. It limited visibility across historical data and hindered rapid detection and coordinated action.	Integrated NVIDIA accelerated computing with the VAST AI Operating System to enable real-time inspection, vector search, automated triage, and unified access to years of telemetry for faster, policy-aware security operations.	Delivered immediate visibility across hot and historical data, reduced analyst burden through agentic automation, accelerated investigations, improved containment speed, and enabled cost-efficient retention of large telemetry volumes for mission-critical environments.
Fluidstack	Fluidstack needed to deliver a reliable, scalable infrastructure for global AI model training. It managed complex customer workloads. It required enterprise-grade performance, security, and data access to support large-scale, distributed AI environments.	Implemented the VAST Data platform to support high-performance AI workloads, enabling scalable storage, multi-protocol access, vector search, and intelligent data processing to manage and accelerate global AI infrastructure deployments.	Improved reliability and performance for large-scale AI training by enabling fast access to distributed data at exabyte scale; enabled secure, future-ready infrastructure supporting global synchronization and efficient training of advanced AI models.

Everest Group® Proprietary & Confidential. © 2026, Everest Global, Inc. | this document has been licensed to **VAST Data**

Stay connected

Dallas (Headquarters)
info@everestgrp.com
+1-214-451-3000

Bangalore
india@everestgrp.com
+91-80-61463500

Delhi
india@everestgrp.com
+91-124-496-1000

London
unitedkingdom@everestgrp.com
+44-207-129-1318

Toronto
canada@everestgrp.com
+1-214-451-3000

Website
everestgrp.com

Blog
everestgrp.com/blog

Follow us on

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at **www.everestgrp.com**.

Notice and disclaimers

Important information. Please read this notice carefully and in its entirety. By accessing Everest Group materials, products or services, you agree to Everest Group's Terms of Use.

Everest Group's Terms of Use, available at www.everestgrp.com/terms-of-use, is hereby incorporated by reference as if fully reproduced herein. Parts of the Terms of Use are shown below for convenience only. Please refer to the link above for the full and official version of the Terms of Use.

Everest Group is not registered as an investment adviser or research analyst with the U.S. Securities and Exchange Commission, the Financial Industry Regulation Authority (FINRA), or any state or foreign (non-U.S.) securities regulatory authority. For the avoidance of doubt, Everest Group is not providing any advice concerning securities as defined by the law or any regulatory entity or an analysis of equity securities as defined by the law or any regulatory entity. All properties, assets, materials, products and/or services (including in relation to gen AI) of Everest Group are provided or made available for access on the basis such is for informational purposes only and provided "AS IS" without any warranty of any kind, whether express, implied, or otherwise, including warranties of completeness, accuracy, reliability, noninfringement, adequacy, merchantability or fitness for a particular purpose. All implied warranties are disclaimed to the extent permitted by law. You understand and expressly agree that you assume the entire risk as to your use and any reliance upon such.

Everest Group is not a legal, tax, financial, or investment adviser, and nothing provided by Everest Group is legal, tax, financial, or investment advice. Nothing Everest Group provides is an offer to sell or a solicitation of an offer to purchase any securities or instruments from any entity. Nothing from Everest Group may be used or relied upon in evaluating the merits of any investment. Do not base any investment decisions, in whole or part, on anything provided by Everest Group.

Everest Group materials, products and/or services represent research opinions or viewpoints, not representations or statements of fact. Accessing, using, or receiving a grant of access to Everest Group materials, products and/or services does not constitute any recommendation by Everest Group to (1) take any action or refrain from taking any action or (2) enter into a particular transaction. Nothing from Everest Group will be relied upon or interpreted as a promise or representation as to past, present, or future performance of a business or a market. The information contained in any Everest Group material, product and/or service is as of the date prepared and Everest Group has no duty or obligation to update or revise the information or documentation.

Everest Group collects data and information from sources it, in its sole discretion, considers reliable. Everest Group may have obtained data or information that appears in its materials, products and/or services from the parties mentioned therein, public sources, or third-party sources, including data and information related to financials, estimates, and/or forecasts. Everest Group is not a certified public accounting firm or an accredited auditor and has not audited financials. Everest Group assumes no responsibility for independently verifying such information.

Companies mentioned in Everest Group materials, products and/or services may be customers of Everest Group or have interacted with Everest Group in some other way, including, without limitation, participating in Everest Group research activities.

Everest Group®
With you on the journey