Endee.io Open Sources its High-Performance Vector Database for Scalable AI

Endee.io launches Endee, an open source vector database delivering fast, accurate, and cost-efficient AI and semantic search at scale.

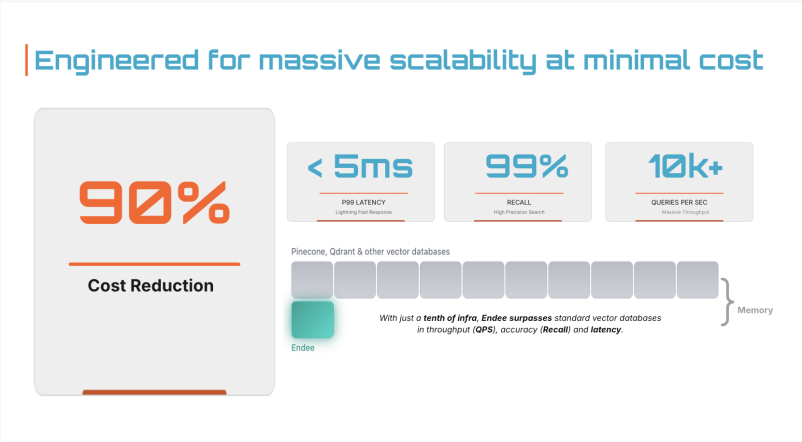

Endee - Next Generation Vector Database For Scalable High-Performance AI

SAN JOSE, CA, UNITED STATES, January 28, 2026 /EINPresswire.com/ -- Endee.io announced the open source release of Endee, a high-performance vector database built for scale, speed, and accuracy. Endee is designed to support modern AI workloads including semantic search, retrieval augmented generation, recommendation systems, and large-scale vector search applications.

As artificial intelligence adoption accelerates, vector databases have become foundational infrastructure for AI-driven applications. Popular platforms such as Pinecone, Qdrant, Milvus, and Weaviate have played an important role in advancing the vector database ecosystem. However, many teams continue to face challenges related to infrastructure cost, memory consumption, and operational complexity as workloads scale.

> Endee rethinks vector DBs for high recall, low latency, and low infra costs. Open sourcing lets teams build faster, secure, and cost-efficient AI systems."
>
> *Vineet Dwivedi*

Endee is built to address these challenges with a focus on high recall, low latency, and efficient infrastructure usage.

A Modern Open Source Vector Database for AI Applications

Endee is an open source vector database engineered to deliver low-latency vector search with high recall while using significantly less infrastructure. Unlike memory-heavy architectures commonly used in vector databases, Endee is optimized to perform efficiently on modest hardware without compromising accuracy.

This makes Endee a strong choice for teams evaluating top open source vector databases for production AI systems where performance and cost efficiency are critical.

Key capabilities of Endee include:

-High performance vector search for AI and semantic retrieval
-High recall with consistently low query latency
-Efficient memory utilization for large-scale vector datasets
-Production-ready architecture designed for extensibility and reliability

Scalable Vector Search with Lower Infrastructure Costs
As vector datasets grow into the millions and billions, infrastructure costs often increase faster than data volume. Endee is designed to scale efficiently, allowing organizations to manage large vector collections without expensive clusters or specialized hardware.

By optimizing indexing, storage, and query execution pipelines, Endee significantly reduces compute and memory overhead. This enables organizations to deploy vector search systems that deliver strong performance while keeping infrastructure and operational costs under control.

Open Source Vector Database Backed by Endee.io
Endee is released as a fully open source vector database, enabling developers, researchers, and enterprises to inspect the codebase, deploy locally, and contribute to its development.

The open source project is led by Endee.io and developed by its parent company, Endee Labs Pvt Ltd, as part of a broader mission to build scalable, secure AI infrastructure. Documentation, examples, and Docker-based setup guides are available to help teams get started quickly.

More information about the project and enterprise offerings is available at https://endee.io/.

Endee Enterprise for Secure and Regulated AI Deployments
In addition to the open source version, Endee.io offers Endee Enterprise, a commercial vector database platform designed for production and regulated environments.
Endee Enterprise supports both serverless deployments and on-premises installations, giving organizations flexibility based on security, compliance, and data residency requirements.

Enterprise features include:

-User and role-based access management
-Queryable encryption to secure sensitive vector and semantic data
-Enterprise-grade security controls and auditability
-High availability, monitoring, and operational tooling
-Performance optimizations for large-scale AI workloads

-Dedicated enterprise support and service level agreements

Founder Quote
“We are extremely proud of how Endee has taken shape and the impact it is already having for teams building AI systems,” said Vineet Dwivedi, Founder of Endee.io. “Our goal with Endee was to rethink vector databases from the ground up so organizations can achieve high recall and low latency without paying the price of massive infrastructure. Open sourcing Endee allows the community to build alongside us and helps teams deploy AI systems that are faster, more secure, and far more cost efficient.”

Common Use Cases
Endee supports a wide range of AI and data-intensive applications, including:

-Semantic search and enterprise search platforms
-Retrieval augmented generation and LLM pipelines
-Recommendation engines
-Knowledge bases and document intelligence systems
-AI agents and real-time inference workflows

Availability
The open source version of Endee is available on GitHub at:

https://github.com/EndeeLabs/endee

Developers can deploy Endee locally using Docker and begin evaluating it as a production-ready open source vector database.

About Endee.io
Endee.io builds high-performance vector database infrastructure for modern AI applications. Through its open source platform and enterprise offerings, Endee.io enables organizations to run scalable, secure, and cost-efficient vector search systems without excessive infrastructure complexity.

Vineet Dwivedi
Endee Labs Pvt Ltd
email us here
Visit us on social media:
LinkedIn
X
Other

try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.