



The AI-Native Vector Database

The easiest way to build and scale AI applications



We're glad to meet you.

- **Founded in 2019** as an AI-native vector database
- **AI-native**, built from the ground up
- **\$50 Series B in April '23**, Backed by Index Ventures, Battery Ventures, NEA, Cortical
- ~65 employees globally (and growing!) majority in Product and Engineering

600+

Paying Customers

4.8

G2 Rating



3M+

Downloads

7K+

Github stars





The future of software is **AI-native**.



The future of software is **AI-native**.

And requires infrastructure that is flexible,
open source, and reliable.



Why Weaviate



Performance & Scale

Scale to Billions of vectors.
Vector compression algorithm.
Native isolated multi-tenancy.



Flexible Deployment Model

Deployable on ANY environment:
SaaS/On-Prem/Hybrid.



End-to-End Focus

Modules allow for easy integration with
the AI landscape.
Simplified developer experience.



Vectors & Objects

Store both Vectors and Objects as
first-class citizens.



Built-In Hybrid Search

Keyword and Semantic together.
Relevance ranking.
No secondary database needed.



Community

3M+ downloads
Support & Documentation



Building and scaling production-ready AI applications isn't easy.

DEVELOPMENT CHALLENGES

Legacy techniques for storing and searching data have limitations.

High barrier of entry for building AI applications. Tool and process silos.

Data privacy concerns around storage and model training.



BUSINESS CONSEQUENCES



High TCO



Poor user experiences



Poor developer experiences



Slowed revenue growth



A flexible, reliable foundation for all of your AI applications.

Semantic search

Image / audio / video search

Recommendations

Question-Answering systems

Chatbots

Agents



Weaviate

AI-Native Development with Weaviate



Hybrid Search

Vectors and objects
Auto-vectorization
Filtering



Retrieval Augmented Generation (RAG)

ML modules
Explainable results
Data privacy



Generative Feedback Loops

Content generation
Automatic tagging
Data enrichment

Developer Enabled Workflows

Extensible platform | Pluggable ML modules | Auto-vectorization | 25+ Integrations

Purpose-Built Architecture

Isolated multi-tenancy | Hybrid search | Filtering | PQ compression

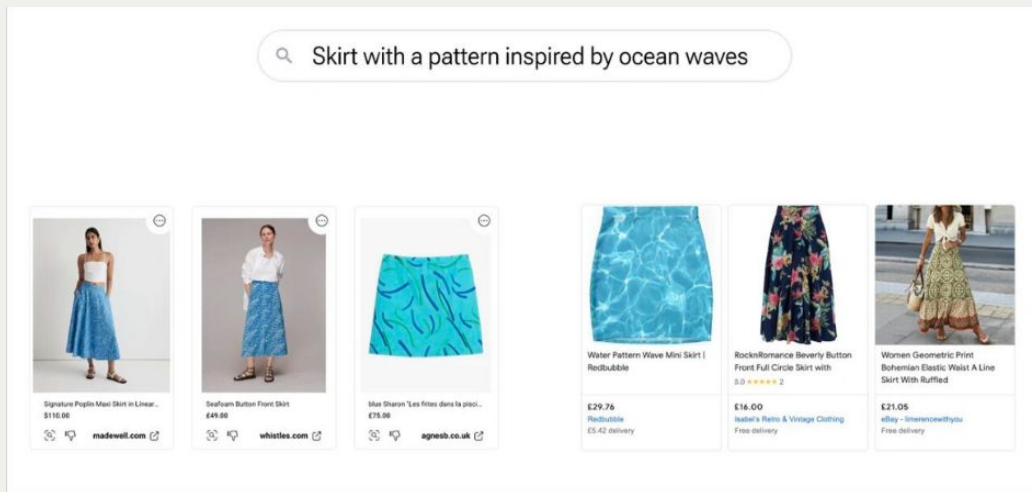
Secure, Flexible Deployment

Open source | Serverless SaaS | Enterprise Dedicated | Bring Your Own Cloud (BYOC)

HYBRID SEARCH

Unlock the power of search to deliver better insights to your customers

- Perform powerful vector and keyword search across all of your data, at scale.
- Simplify development with pre-built modules for vectorization and LLMs.
- Deliver results with the speed and accuracy users expect, efficiently.





RETRIEVAL AUGMENTED GENERATION (RAG)

Build secure, stateful, explainable AI applications with less hassle

- Use proprietary data for LLM-based apps, without sacrificing privacy.
- Easily plug in different ML frameworks and models to find what suits your use case.
- Scale with a reliable, flexible, open source platform that adapts to your needs.

What's the CAGR from 2020 to 2022?
Apple Form 10-k.pdf

To calculate the CAGR (Compound Annual Growth Rate) from 2020 to 2022, we need the net sales figures for those years. From the text, we have:

- 2020 net sales: \$274,515 million
- 2022 net sales: \$394,328 million

The formula for CAGR is:

$$\text{CAGR} = (\text{Ending Value} / \text{Beginning Value})^{(1/ \text{Number of Years})} - 1$$

Plugging in the values:

$$\text{CAGR} = (394,328 / 274,515)^{(1/2)} - 1$$

CAGR= 1.1991
CAGR= 0.199 or 19.9%

The CAGR from 2020 to 2022 is approximately 19.9%.

Page 25 / 80

Active File | All Files | Advanced

What's the CAGR from 2020 to 2022?

Japan net sales decreased during 2022 compared to 2021 due to the weakness of the yen relative to the U.S. dollar.

Rest of Asia Pacific
Rest of Asia Pacific net sales increased during 2022 compared to 2021 due primarily to higher net sales of iPhone, Mac and Services. The weakness in foreign currencies relative to the U.S. dollar had an unfavorable year-over-year impact on Rest of Asia Pacific net sales during 2022.



“Weaviate included model serving and multi-tenant implementation which have helped us quickly prototype and build our vector search at Stack.”

Constantin Kokkinos
Staff Data Platform
Engineer



“Within a couple of weeks we had a novel, production-ready AI-powered search engine. Easy to set-up, good documentation, and responsive support. Exactly what we needed.”

David Wood
CTO



“Building with AWS and Weaviate allows us to innovate at the speed customers need while keeping accuracy, data privacy, and explainability top of mind.”

Anant Bhardwaj
CEO

58M

Questions and answers

60M

Product SKUs

Unlimited

Documents



Ask like a human: Implementing semantic search on Stack Overflow

Semantic search allows users to search using natural language instead of a rigid syntax of keyword manipulation.

Search has always been core to the Stack Overflow experience. It was one of the first features we built for users, and since the early days of Stack Overflow, most of our visitors have come to us through search engines. With search engine referrals consistently comprising over 90% of our site traffic, we felt for many years as though this was a solved problem.







Building an AI-Powered Shopping Copilot with Weaviate

UK-based startup Moonsift is harnessing the power of AI—using machine learning models and Weaviate’s vector database—to help online shoppers discover the products they love. [Moonsift](#) offers an ecommerce browser extension for users to curate shoppable boards with products from across the internet.

Skirt with a pattern inspired by ocean waves



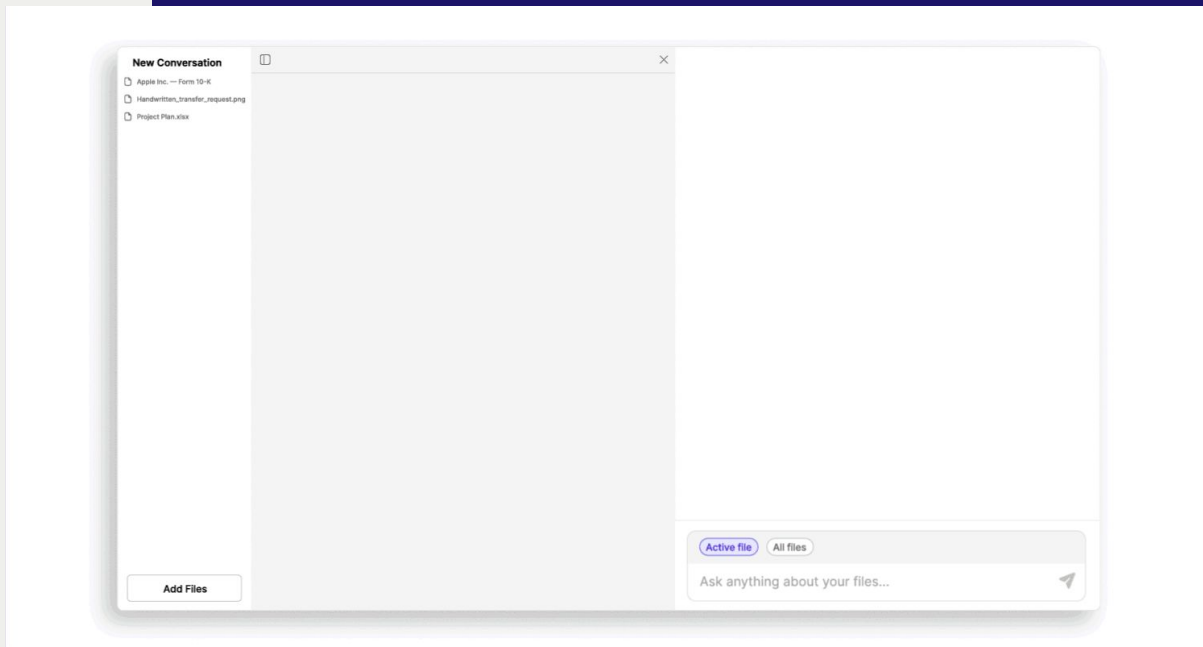
vs.





Engage in a conversation with your content

[AI Hub Converse](#) lets you create a conversation with your content to extract insights, analyze, [summarize](#), [translate](#), reason, or even generate new content. Instabase creates a digital representation of your content, from handwritten notes to spreadsheets, that gives GPT-4 a newfound understanding of structure, style, and meaning.






Why build with Weaviate

- **Pluggable ML models**
- **Built-in hybrid search**
- **Developer experience & ecosystem**
- **Unique multi-tenant architecture**
- **Flexible, secure deployment**



Integration to the AI ecosystem



LangChain



haystack
by deepset



Gemini



PaLM 2



Amazon
SageMaker

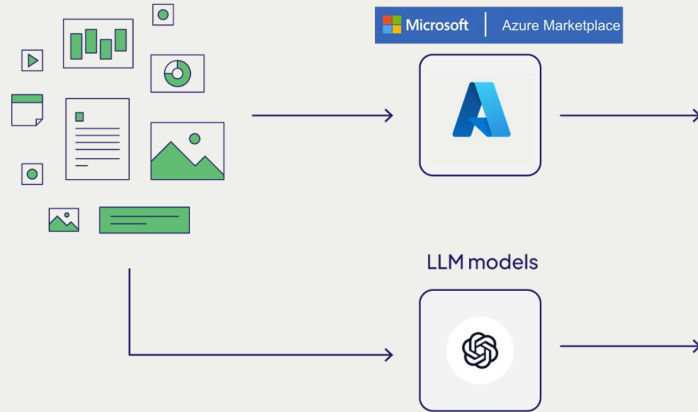


Security, privacy, and reliability

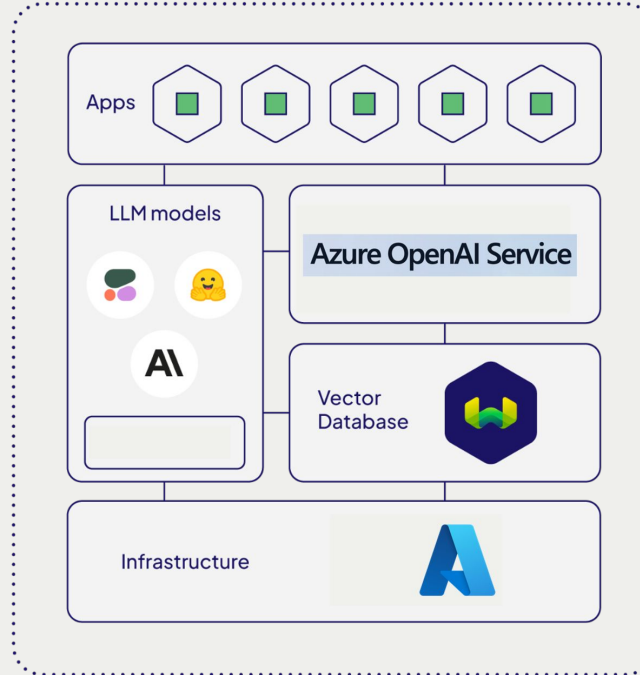
Secure Deployment	End-to-end Encryption	User Authentication	Native Multi Tenancy
Run in a dedicated tenant or in your own VPC.	Data is fully encrypted in transit and at rest.	Connect to your identity provider to rely on built-in mechanisms.	Isolate data with industry-leading tenant management.
Advanced Firewall Protection	Active Monitoring	Multi Availability Zones	Automated Backups
Take extra precaution against malicious threats.	We proactively monitor our systems and are on standby 24/7 for incidents.	High availability and resilience across 3 zones.	Rest easy with daily backups and snapshots.

Architecture Diagram

Any data type ranging from text, audio, image, depth, healthmap, more...



Optional, direct integration with embedding and generative AI providers



One infra,
One data store
No secondary
database
needed.



Thank you

azure@weaviate.io



Connect with us!



weaviate.io



[weaviate/weaviate](https://github.com/weaviate/weaviate)



[weaviate_io](https://twitter.com/weaviate_io)