QDRANT SOLUTIONS GMBH

Qdrant

Enterprise-Ready, Massive-Scale Vector Search Technology for the Next AI Generation

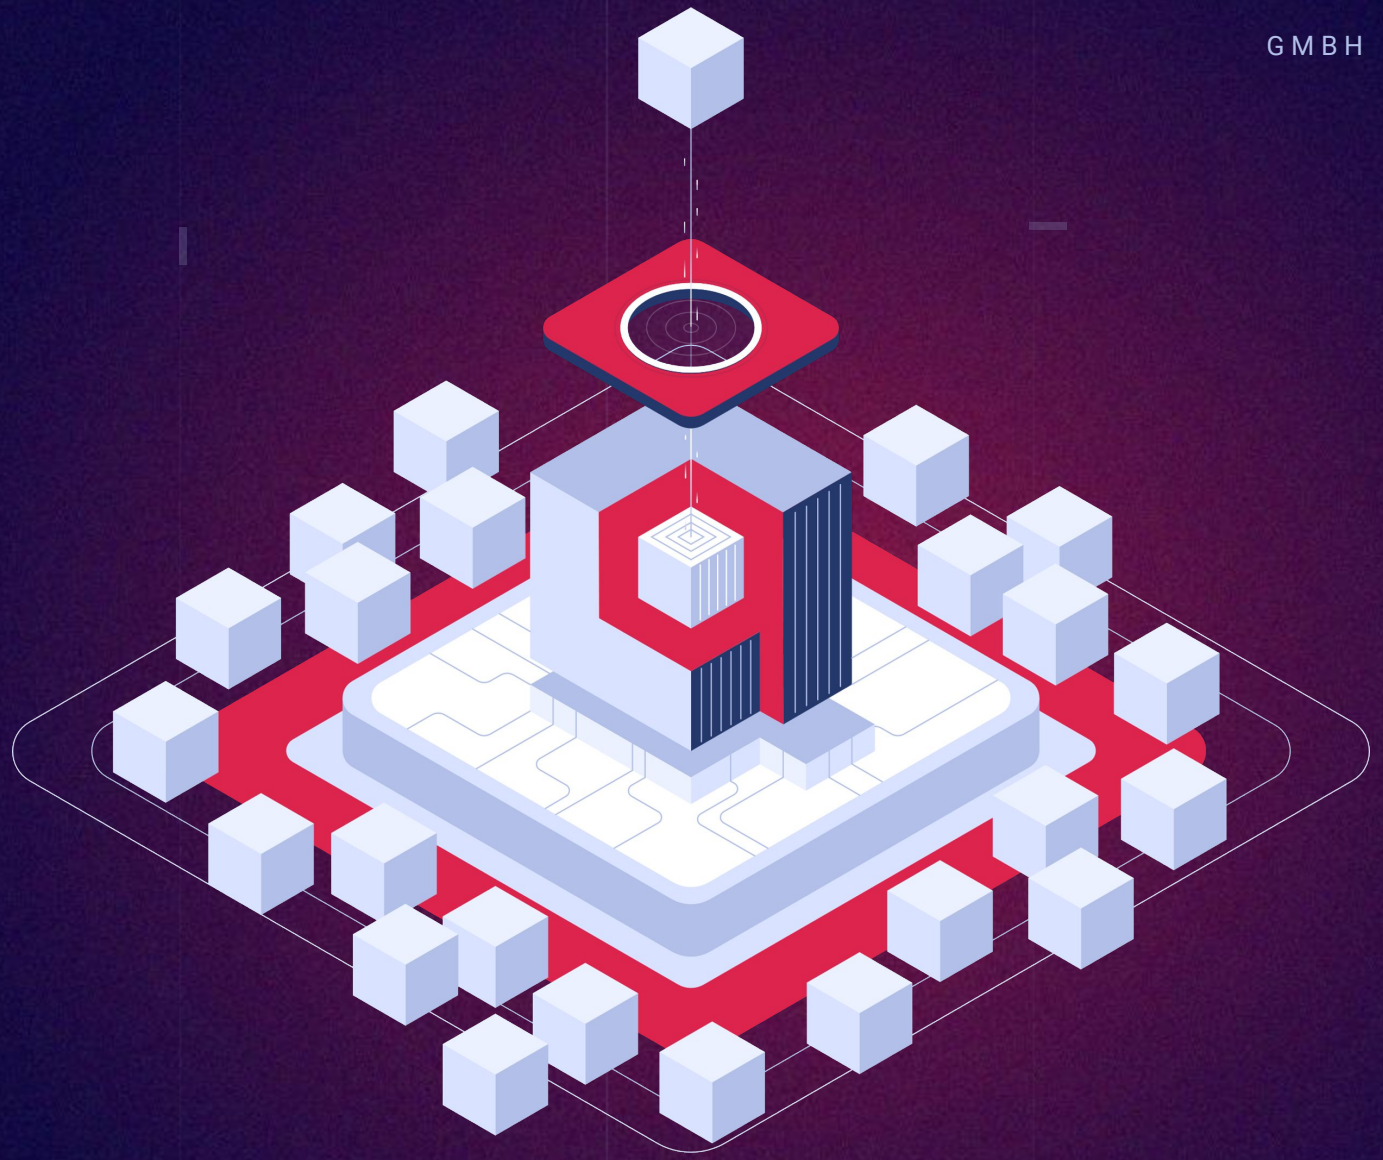

PRODUCT DECK

CONFIDENTIAL

Qdrant at a glance

Overview

CHARACTERISTICS

- **Best available performance, blazing fast and accurate.**
- **Rich feature set for resource optimization.**
- **High-scale, distributed, durable and cloud-native.**

HIGHLIGHTS

- Fastest, feature-rich, advanced open-source vector database.
- Available as open-source, managed cloud or on-premise deployment
- Integrated with popular frameworks like jina.ai, LLaMa Index, LangChain, Haystack, etc.
- Partnering with LLM providers like Cohere AI, Aleph Alpha, OpenAI
- Used and trusted by top innovative companies worldwide

TRACTION

Top

open-source vector search database on GitHub, >14K Stars

> 10 K

OS Adopters Worldwide

> 5 M

Downloads

Demand for AI outpaces infrastructure

The global datasphere will grow to 163 zettabytes by 2025, and about **80% of that will be unstructured**

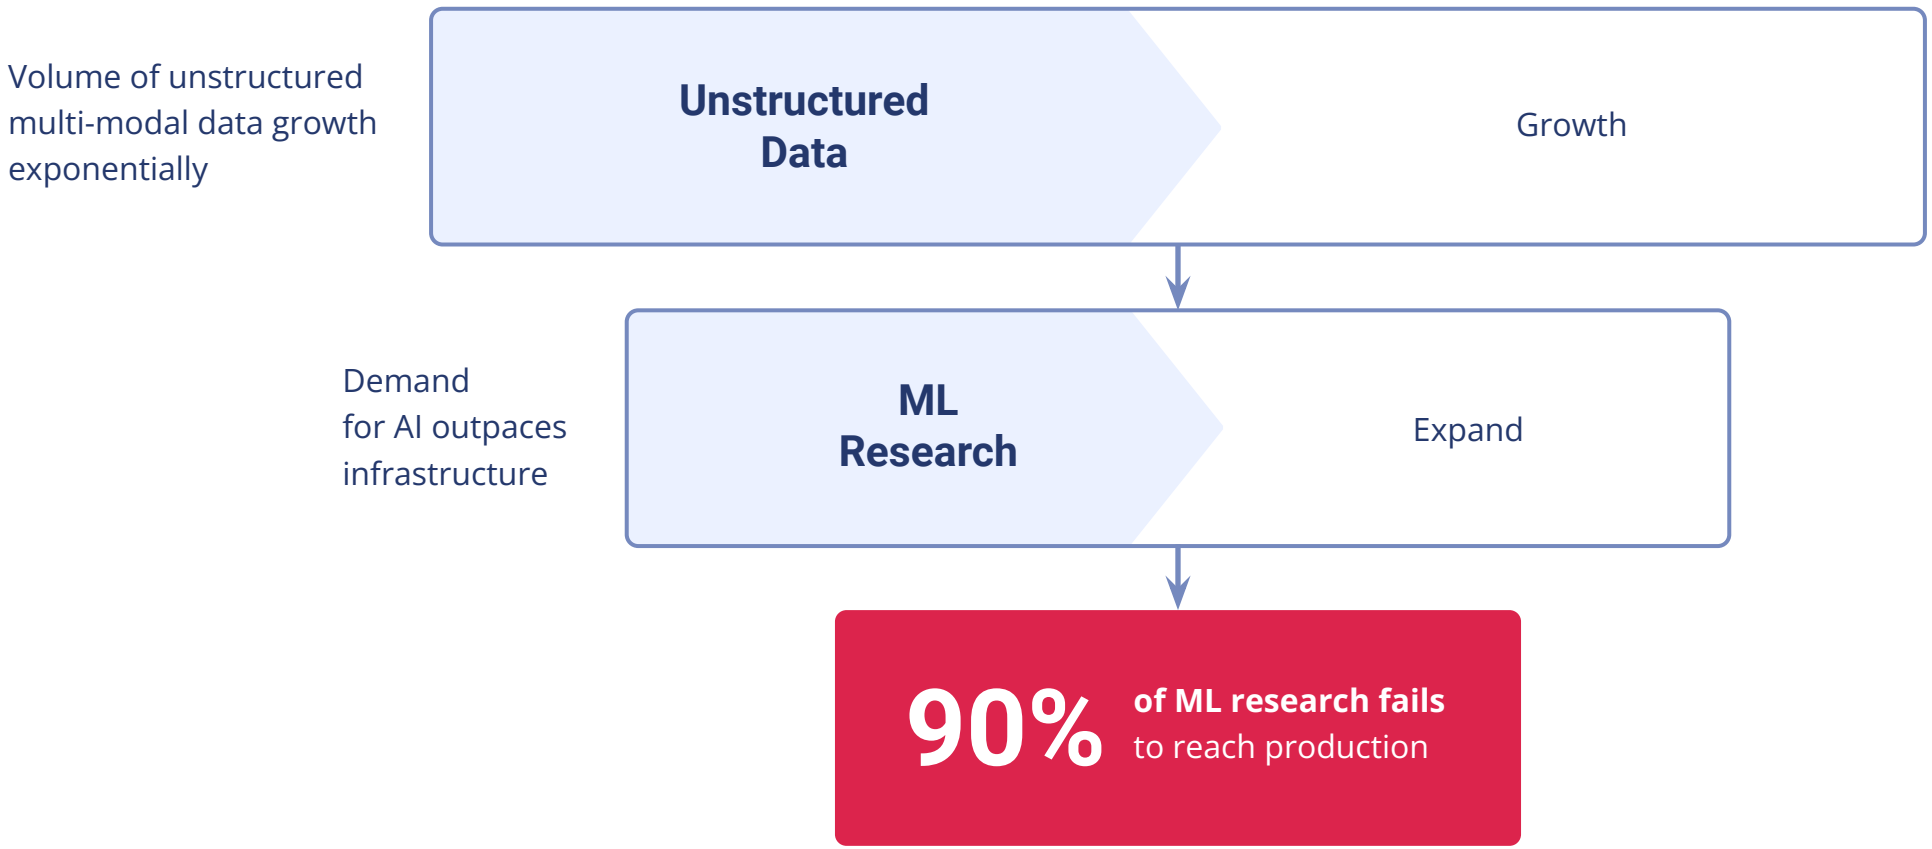

Vector Search

An essential part of the AI Transformation

drant

INDUSTRIES

HR-Tech | Ad-Tech | Online Dating | Gig Economy | E-Commerce | Law-Tech | Fashion | Ed-Tech

Media & News | Biometrics | Med-Tech | Anti-fraud | Agriculture | Manufacturing | Streaming Services | Marketplaces

SEARCH SYSTEMS

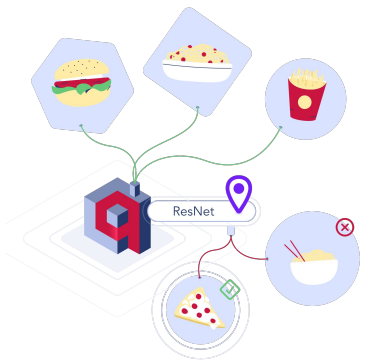

MATCHING ENGINES

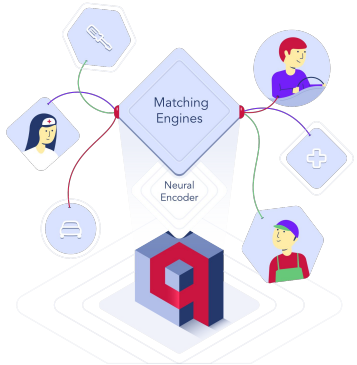

ANOMALY DETECTION

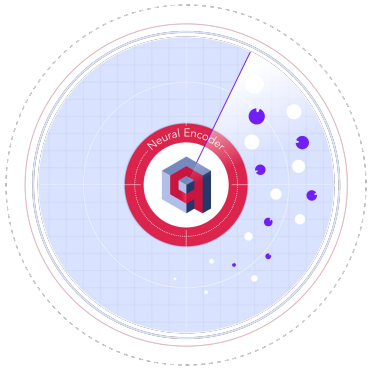

RECOMMENDATIONS

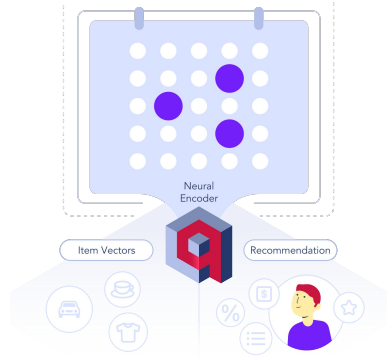

NN Encoders + Vector Database

Emerging use cases

Fundamental building block of the new AI Stack

LLM Memory

Copilot for X

Neural Frameworks

deepset txtai jina LangChain

Vector Database

qdrant

ChatGPT/LLMs

cohere OpenAI Transformers

Qdrant – the fastest Vector Database

BY SEARCH SPEED

Up to 10 times faster
according to benchmarks

BY ADOPTION

>10000 Production
Deployments

BY DEVELOPMENT SPEED

Get started **in minutes**

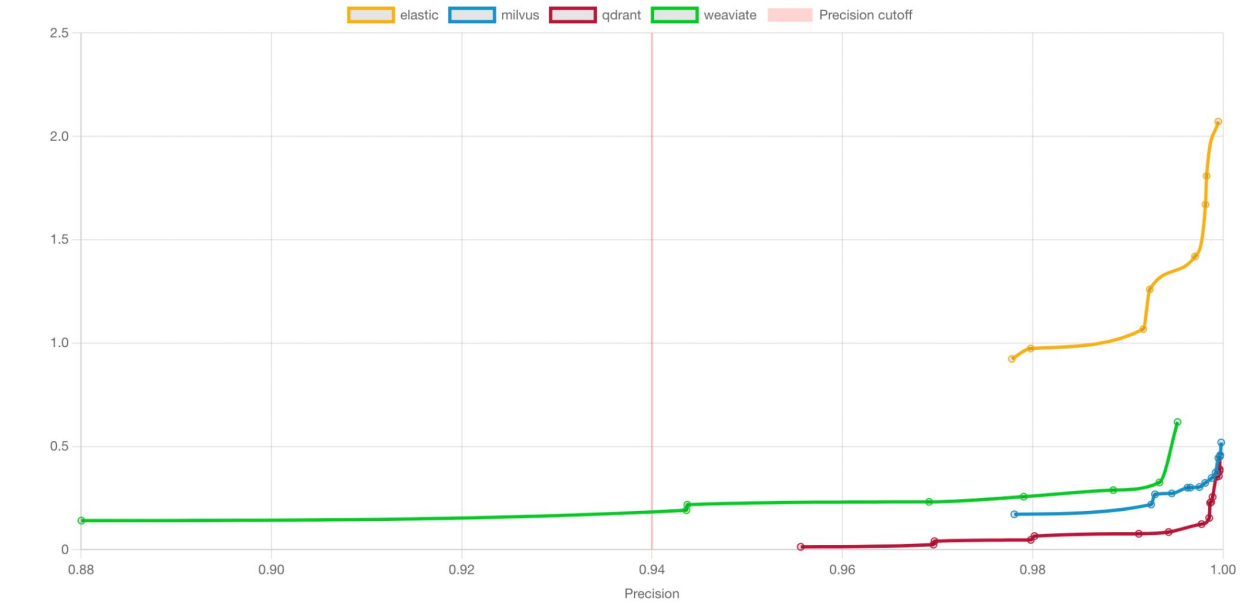

Engine Architecture

Well-thought from Hardware Optimization to Distribution with Raft

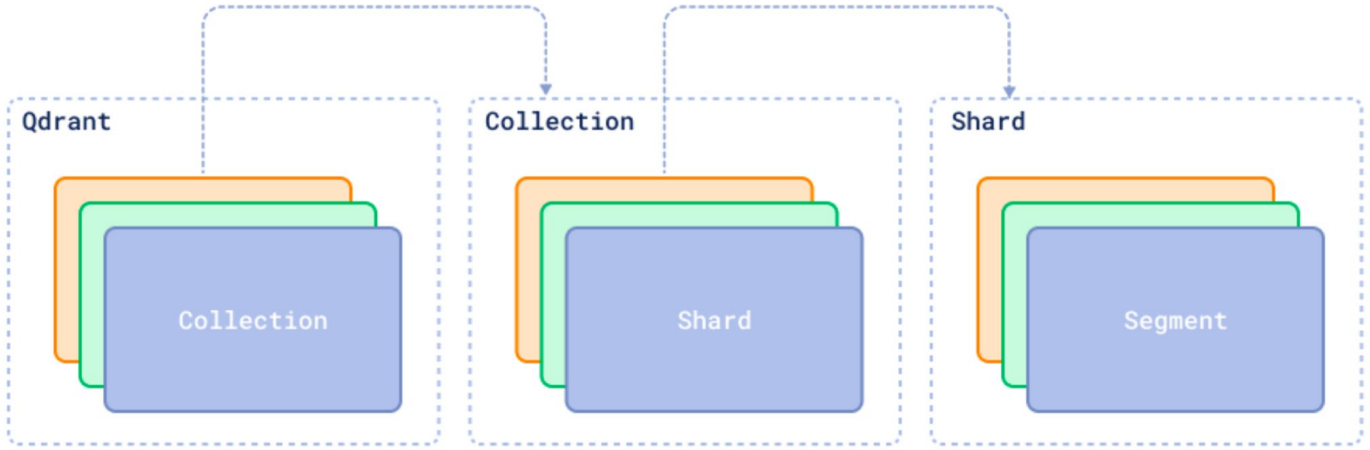

Designed for Performance & High Availability

- Implemented from scratch entirely in Rust
- Flexible architecture with segments on storage level
- Horizontal scaling out with sharding and distributed deployment.
- Replication for HA and throughput gain

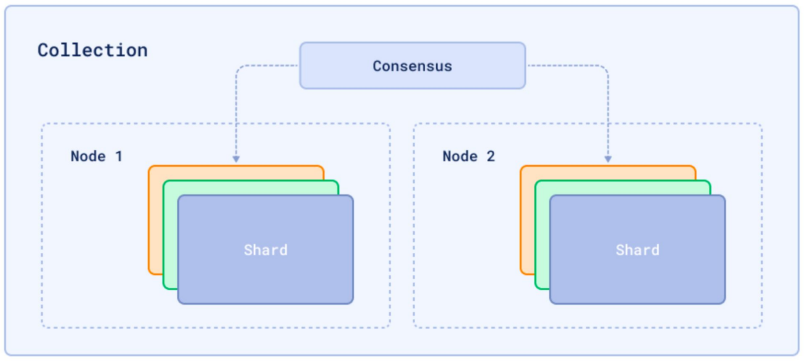

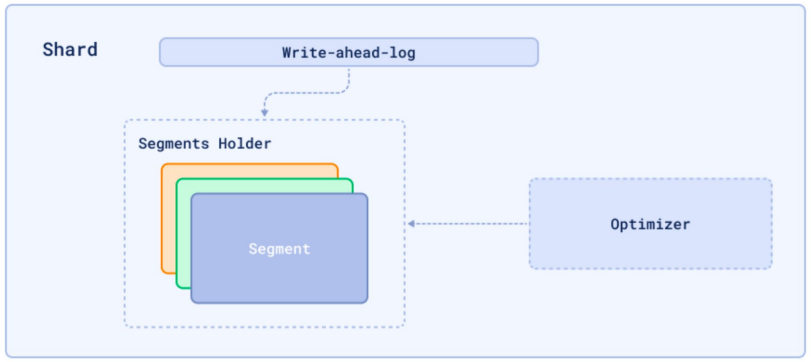

Superb Features

Flexibility in cost and performance setup options

- **Advanced one-stage Filtering**
  Store any kind of data including geolocation along with vectors and filter search results.
- **Scalar, Product, and unique Binary Quantization**
  Up to 64x Memory usage reduction and x40 search speed up.
- **Memory Maps and IO Uring**
  Effective on-disk storage options and low level hardware optimization.
- **Zero-downtime** update of vectors in case of model switch (blue-green deployment)
- **Advanced multi-tenancy** support
- **Read/Write segregation** and dynamic read scaling
- **TTL and large bulk operations** without affect on retrieval speed
- **Payload-based sharding**

- **Unstructured Data Visualization**
  Vectors visualization dashboard using t-SNE dimensionality reduction algorithm.

Product Overview

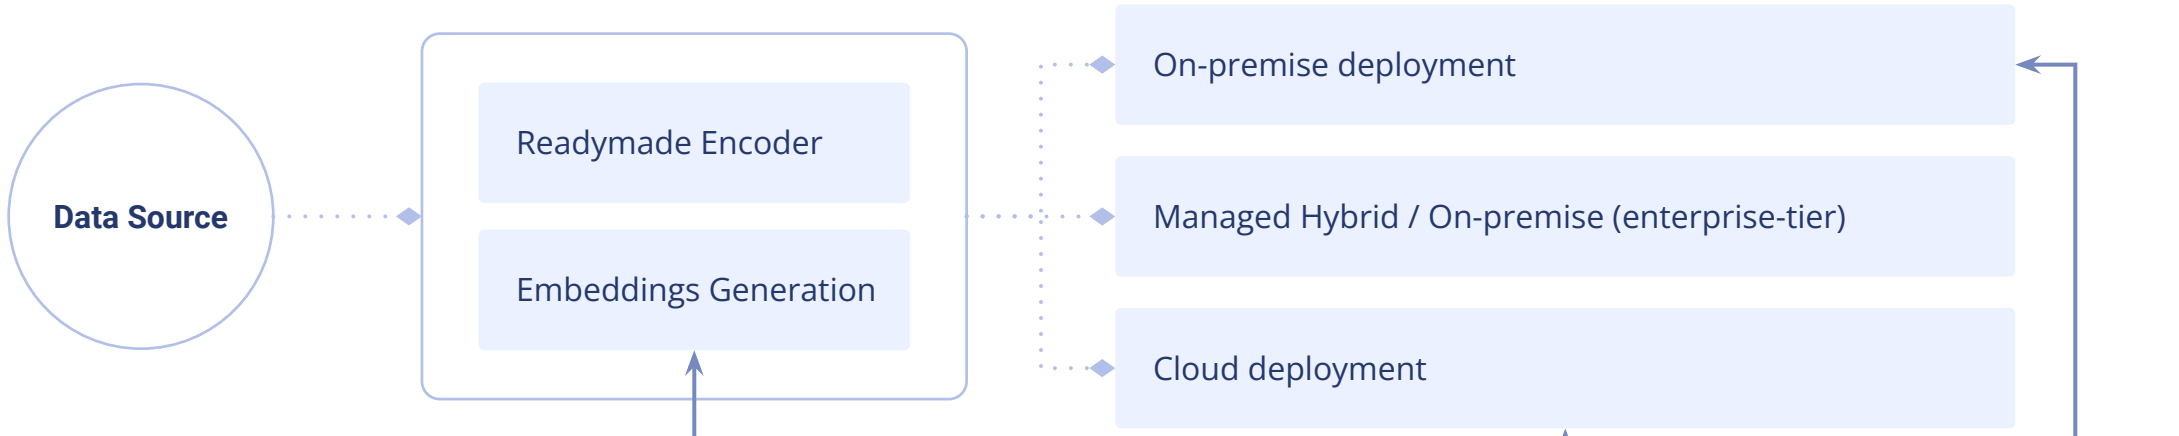

FastEmbed

Embeddings generation library

- Open-Source Models
- Accuracy with high Recall
- Fast even without GPU

Managed Cloud SaaS and On-prem

- Access management
- Monitoring & Alerts
- Enterprise-scale features

Open-Source Vector Database

- REST & gRPC service
- Extended filtering support
- Horizontally scalable
- Written in Rust

Feedback

What Qdrant users have to say

Hooman Sedghamiz

Director AI /ML, **Bayer**

Vectorstores are definitely here to stay, the objects in the world around us from image, sound, video and text become easily universal and searchable thanks to the embedding models and vectorstores.
I personally **recommend Qdrant**. we have been using it for awhile and **couldn't be happier**.

Andrew Rove

Principal Engineer, **GLG**

We have been using Qdrant in production now for over 6 months to store vectors for cosine similarity search and it is **way more stable and faster** than our old ElasticSearch vector index. No merging segments, no red indexes at random times. It just works and was super easy to deploy via docker to our cluster.

Christopher Prohm

Data Scientist & Topic Lead, **Volkswagen**

My feedback is: **Qdrant is a great** piece of tech!

Hamza Farooq

Sr. Researcher, **Google**, Lecturer, **Stanford**

I've been a really **big fan of Qdrant** and teaching a course on LLMs using qdrant for storing embeddings and indexing.

Olaf Górski

Sr. AI Engineer, **SumUp**

I was evaluating different vector store for use in our upcoming launch of a new production feature. TLDR: **qdrant won** for our use case in certain aspects, hence. I made the decision to roll with it.
The **ease of use** is & getting started is **amazing!**

Prashanth Rao

Sr. AI Engineer, **Royal Bank of Canada**

I absolutely **love using Qdrant**, and I really hope you secure a huge round of Series A/B funding in the years to come! Looking forward to evangalizing more amazing Rust-based software like yours, so thank you for building something so cool!

More feedback

What Qdrant users have to say

Adrien Biarnes

Senior DS @ **Dailymotion**

...So I looked into it and found that VertexAI had exactly this service called Matching Engine. So we tried it out but we also found Qdrant.
So we benchmarked the 2 **and it was a no-brainer.** It cost us a big pile of money to create the indexes with Matching Engine... Also, the Matching Engine API is very poor compared to Qdrant. Considering the price of the service it almost feels like being ripped off.

Arbaaz Khan

ML Infrastructure @ **Doordash**

Our vector search is mostly around similarity search with metadata filtering. There are several different types of application use cases that depend on this central concept and would have disparate scale requirements ranging from 10M–500M vectors. Our benchmark testing gave remarkable performance vs milvus. Already starting to **see Qdrant emerge as the winner.**

Alex Webb

Director of Engineering @ **CB Insights**

We looked at all the big options out there right now for vector databases, with our focus on ease of use, performance, pricing, and communication. **Qdrant came out on top in each category**... ultimately, it wasn't much of a contest.

Mike Gugel

Data engineer @ **Bauhaus**

We're trying Qdrant lately in a PoC-like project dealing with mass similar images search as the vector-db for storing the embeddings and pretty happy with it, thank you for creating such an awesome piece of software, it's blazingly fast, super dev-ops friendly, very convenient to use, **totally a big like!**

Jon Tse

Senior director of R&D @ **JCV (by Softbank)**

Qdrant is the best Vector Similarity Search Engine ever, as **every other VSS broke** during our our high load and performance testing. If we find any issues, the Qdrant team fixes the fastest and improves on things beyond!

Thank You

sales@qdrant.com