

aws marketplace

Qdrant

Reviews, tips, and advice from real users



Powered by  PeerSpot

Contents

Product Recap..... 3 - 4

Valuable Features..... 5 - 11

Other Solutions Considered..... 12 - 14

ROI..... 15 - 17

Use Case..... 18 - 22

Setup..... 23 - 24

Customer Service and Support..... 25 - 26

Other Advice..... 27 - 31

Trends..... 32 - 33

About PeerSpot..... 34 - 35

Product Recap



Qdrant

Qdrant Recap

Qdrant is a powerful tool for efficiently organizing and searching large volumes of data. It is particularly useful for tasks such as data indexing, similarity search, and recommendation systems.

With fast and accurate results, it is suitable for various applications including e-commerce, content management, and data analysis. Users appreciate Qdrant's efficient search capabilities, high performance, and ease of use.

Its quick and accurate retrieval of relevant information allows for easy navigation and analysis of large datasets.

The intuitive interface and straightforward setup process make it accessible to users with varying levels of technical expertise.

Valuable Features

Excerpts from real customer reviews on PeerSpot:



“Qdrant has positively impacted my organization by significantly improving our RAG system's latency, cost-efficiency, and accuracy.”



Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees



“Qdrant has positively impacted my organization by consuming much less time than building systems through coding.”



Chirag Morajkar

Lead Ai Tech And Tech Automation Engineer at a individual & family service with 11-50 employees



“We saw a clear return on investment from Qdrant, particularly in the engineering time saved and the empowerment of team members to handle self-service tasks instead of reducing headcount.”



Amrit Dash

Automation Engineer at a educational organization with 11-50 employees

- ✓ “Qdrant is an excellent vector database that anyone would want to use with RAG AI.”



Validated user

Product Engineer at a tech vendor with 11-50 employees

- ✓ “Qdrant has reduced our response time to less than one second for our 128 KB token sizes, and we are satisfied with that performance.”



China Venkanna Varma Ponnamanda

Co Founder & CEO at SaYukth Private Limited

- ✓ “Qdrant is one of the best vector databases out there that is also open source.”



Adib M

Founder at a tech vendor with 11-50 employees

- ✓ “Qdrant has positively impacted my organization mainly by reducing usage costs for AI because in the older system, in order to get the context, we were just dumping everything into the system, which would inflate the cost of the API, whatever AI we were using, such as OpenAI or Claude.”



Mir Shahzad Mubeen

CTO at Honeycomb AI

What users had to say about valuable features:

“The best features Qdrant offers are the ability to talk to the AI agents and transform the responses. When I use it in my N8n workflow, it makes my RAG workflow respond very well and then I can achieve all my desires on the project.

“Regarding Qdrant's AI capabilities, its accuracy and reliability of output are very accurate. I have used it several times and I enjoy using it a lot..”

Femi Adeyemi

Senior Team Lead Systems Administrator at a comms service provider with 51-200 employees

[Read full review](#) 

“The best features Qdrant offers are that it is easier to set up, and there is nothing too complex about it. The dashboard visualization is also quite good. Qdrant supports high-dimension vectors and cosine similarity, which any vector database should have, and it is pretty fast.

We do not directly use the visualization feature in our work, but I do check it occasionally to see how many data fields we have. It might help in determining how much we are using and what the data fields are. In the workflow or pipeline itself, it is not used for any case at this time..”

Mir Shahzad Mubeen

CTO at Honeycomb AI

[Read full review](#) 

“Qdrant can handle high-dimensional vectors quite well. Other options, such as Chroma or Faiss, are mostly for POCs. Qdrant is a proper one, and there are others also, such as Pinecone. However, Pinecone or AWS Search are all paid options and services. There are advantages of paid services as well as advantages of open source.

“Qdrant also has a paid version where the vector database is given as a service. Your collections are created on the cloud, so the pressure on your infrastructure is less. If you keep creating the vector database collections, it will produce a strain on your machine. That is why you need a service, and I have not used the service, but given a requirement, we might end up using it. By service, I mean the vector database gets created on the cloud itself..”

Uzair Faruqi

Head Of Data Science at Mjunction Services

[Read full review](#) 

“I mostly value storing all the embeddings, and Qdrant also has really good support for different types of vector storage. For example, HNSW and IVF are pretty good options. Qdrant is the best and most batteries-included framework or vector database out there, which is also open source. Pinecone is good but it is closed source. Qdrant is my next best option for building something scalable.

“The biggest advantage that we have is that we can self-host the entire thing. By self-hosting this, we can allow ourselves to explore the technology a bit more. We also have the confidence and trust about the technology being open source, so even if the company collapses or something happens, we can always go back and continue building the technology. It is mostly the trust and the flexibility to add our own features or build on top of it.

“Qdrant is pretty accurate, and one of the best features is that everything can be tweaked via parameters. I can increase the M, I can increase the number of node connectivity, and so on, via Qdrant's Python or any sort of SDK. I can always have control over accuracy versus speed, which is amazing for Qdrant..”

Adib M

[Read full review](#) 

Founder at a tech vendor with 11-50 employees

“Beyond standard retrieval, Qdrant's payload indexing is critical for our multi-query setup. It allows us to pre-filter dense vectors by document type, such as instructions versus definition bases, before running reciprocal rank fusion, which drastically reduces search latencies and eliminates out-of-context hallucinations from our AI responses.

“In my experience, the best features Qdrant offers include advanced payload filtering, HNSW index efficiency and high-dimensional support, hybrid search and multi-query integration, as well as managed cloud reliability and REST gRPC APIs.

“Two major features stand out from an architecture and operations perspective. One is one-stage filtering during HNSW traversal, and the second is Rust-based performance and resource efficiency.

“Qdrant has positively impacted my organization by significantly improving our RAG system's latency, cost-efficiency, and accuracy. For example, we experience faster response times, lower API and compute costs, and zero hallucination on document scope.

“By shifting our filtering and multi-query setup into Qdrant, we see immediate improvements in system performance and operational costs. Our raw vector retrieval and payload filtering drop down to under 15 milliseconds. This makes a substantial difference in keeping our total end-to-end pipeline fast, even when running multi-query fusion loops with LangGraph. Metadata pre-filtering by document type means we only pass precise, highly relevant text nodes into our prompt. By trimming out unnecessary context and tightening our top retrieved documents, we reduce our input token consumption on Azure OpenAI GPT-4o by approximately 30 to 40 percent. Beyond the cost savings, restricting retrieval strictly to verified document categories effectively eliminates out-of-domain hallucinations..”

Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees

[Read full review](#) 

“Qdrant's most important feature is that it is very fast because it tends to give the results quickly. It is also very compatible with Python because Qdrant has first-class Python support. I can quickly install it on my machines with a PIP install Qdrant client, allowing me to work smoothly.

“The other important feature is its cloud capability. Qdrant offers self-hosting, and I am using it from Docker, which is a significant advantage. It has high performance and is good for flexible deployment. Another main advantage is that it is very cost-efficient since it is open source.

“The biggest importance I have seen in Qdrant is its high performance. It is a Rust-based engine with very low latency and HNSW indexing. High performance is a big plus when using Qdrant, especially when I need to fetch data through the vector database for RAG AI, as it tends to provide outputs quickly. Since I use Python, it is very fast and API friendly, and that is one of the best things I have experienced.

“Qdrant works well with LLM web frameworks, especially with native support in LangChain, where the integration is pretty easy from the LangChain perspective.

“Qdrant has improved my efficiency and is really easy to integrate. It is self-hosted, so integration is straightforward. I believe it offers many advantages compared to what I have used alternatively before.

“Qdrant has saved a lot of time and has improved efficiency and reduced costs since it is open source and very cost-effective. Although I cannot specify a particular number, I estimate that I have saved about 50 to 60 percent of the time by using Qdrant. There is no extra charge since it is open source, which is a significant benefit..”

Validated user[Read full review](#) 

Product Engineer at a tech vendor with 11-50 employees

Other Solutions Considered

“Though the leader and the best vector database in the field is considered to be Pinecone, which is a paid cloud service, that I would consider to be the benchmark in the industry. Qdrant is good, but Pinecone is also much more costly..”

Uzair Faruqi

Head Of Data Science at Mjunction Services

[Read full review](#) 

“Prior to selecting Qdrant, I evaluated several other popular vector database options, primarily Pinecone, ChromaDB, and PGVector. Ultimately, Qdrant gives us the best balance of open-source transparency, Rust-baked latency performance, managed cloud options, and superior metadata payload filtering..”

Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees

[Read full review](#) 

“I did evaluate other options before choosing Qdrant. I looked at Chroma and Pinecone. Chroma was probably my second best choice, but I did not really prefer Pinecone too much..”

Mir Shahzad Mubeen

CTO at Honeycomb AI

[Read full review](#) 

“Unfortunately, not by much because we already did a lot of research beforehand and picked out Qdrant without spending a lot of time exploring multiple different vector databases. I did look into other vector databases, but I did not use any of them at that scale where we could determine how much money was saved. I do not have a lot of metrics for comparison, unfortunately. However, I did conduct some initial benchmarking and Chroma is quite fast. Chroma has issues with scaling. Qdrant used to be in the second or third position and was, in general, quite a good vector database..”

Adib M

Founder at a tech vendor with 11-50 employees

[Read full review](#) 

“I initially prototyped with ChromaDB because it was easy to set up locally during early development. However, as our assistant matured into a complex, multi-query RAG application with LangGraph, Chroma began to hit limits around metadata filtering speed, payload management, and production-grade scaling. We switched to Qdrant because of its single-stage HNSW payload filtering, sub-15 millisecond query latencies, and seamless managed cloud offering, which gave us the exact performance, reliability, and API flexibility we needed for production..”

Pawel Cislo

[Read full review](#) 

MLOps Engineer at a tech services company with 501-1,000 employees

“Before using Qdrant, I tried several alternatives such as Pinecone, which also manages vector databases but is more expensive at scale and closed source, and ChromaDB, which is suitable for prototyping but does not perform well for production scaling. In contrast, Qdrant is fast, offers strong filtering, allows easy self-hosting, provides good cloud options, and is simpler to use than Milvus, Weaviate, or Pinecone.

“I evaluated several other options, including Pinecone, Weaviate, Milvus, and ChromaDB. They each have downsides, particularly in production scaling, complexity of architecture, or being more resource-intensive compared to Qdrant..”

Validated user

[Read full review](#) 

Product Engineer at a tech vendor with 11-50 employees

ROI

Real user quotes about their ROI:

“I have seen a significant return on investment from using Qdrant because it is very easy to integrate and highly efficient, saving a lot of time in my day-to-day operations, which ultimately saves money as well. It is productive as I can do more work in less time..”

Validated user

Product Engineer at a tech vendor with 11-50 employees

[Read full review](#) 

“I have metrics on my API costs which have definitely shown a thirty to forty percent drop because of Qdrant. Overall, I would say pipeline processing efficiency has increased, and there are definitely fewer issues with using an AI agent with our client because context is now loaded rather than just dumped..”

Mir Shahzad Mubeen

CTO at Honeycomb AI

[Read full review](#) 

“I see a clear return on investment across cost savings, response speed, and development efficiency. By utilizing Qdrant's payload filtering to restrict searches to specific document categories, we reduce the volume of irrelevant context sent to Azure OpenAI. This lowers our LLM input token consumption by roughly 30 to 40 percent, translating directly into lower monthly OpenAI API bills..”

Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees

[Read full review](#) 

“There is definitely a return on investment with Qdrant. Although I am on a free tier and therefore have zero investment, the results obtained have been excellent for a vector database. The time saved is substantial, with nearly three weeks or more for projects deployed with Qdrant Cloud in no-code platforms. The accuracy achieved through HNSW searching compared to others is notable..”

Chirag Morajkar

Lead Ai Tech And Tech Automation Engineer at a individual & family service with 11-50 employees

[Read full review](#) 

“We saw a clear return on investment from Qdrant, particularly in the engineering time saved and the empowerment of team members to handle self-service tasks instead of reducing headcount. The most significant ROI stemmed from improved ticket resolution time, with issues being solved rapidly as the team became more autonomous. Thanks to Qdrant's open-source nature, our initial licensing and setup costs were nearly zero, allowing for swift testing and launch of our RAG prototype. Although maintaining dual databases created some complications, switching to Supabase Vector preserved those productivity enhancements while significantly alleviating DevOps time and reducing costs..”

Amrit Dash

Automation Engineer at a educational organization with 11-50 employees

[Read full review](#) 

Use Case

“My main use case for Qdrant is local AI development. A specific example of how I'm using Qdrant in my development work is that I use it for my RAG development using N8n. In my drug development work, I use Qdrant on my N8n workflow and it has helped me with my local AI integration with the servers..”

Femi Adeyemi

Senior Team Lead Systems Administrator at a comms service provider with 51-200 employees

[Read full review](#) 

“We have a full-fledged RAG system using Qdrant Vector Database, and that is how it has benefited us. For example, we have implemented a techno-commercial evaluator using that, and it is in production. Qdrant is a good database and a good vector database. They have an on-premise option also, using a Dockerized system, and you can use their cloud services..”

Uzair Faruqi

Head Of Data Science at Mjunction Services

[Read full review](#) 

“Qdrant's main use case is to fetch data through the vector database to populate my RAG AI. When I input data into the RAG AI, Qdrant stores the text, images, or other data as embeddings, or vectors, and it helps the RAG system retrieve the most semantically relevant information based on user input. The most common data types could be text, images, audio, or any data that can be converted into a vector. Once I put the user input, it provides an output through those embeddings.

“That is the main workflow I've been using Qdrant for the past year. I've also been using it in LangChain because it works really well with LangChain. Beyond that, I use Qdrant to upload snapshots as well..”

Validated user

[Read full review](#) 

Product Engineer at a tech vendor with 11-50 employees

“My main use case for Qdrant is as a vector database. As a vector database, I use Qdrant in a sequence and pipeline automation software that needs to look up information about certain items, each of which has loads of information associated with it. In order to make it easy for the AI to have the context in hand, we deployed a vector database with a very high 1024 dimensions so that the AI can read context and analyze as it is supposed to, rather than just dumping everything into the prompt.

I decided to use Qdrant for this pipeline automation because I wanted a vector database that supports high-dimension vectors like the 1024 dimensions we are using. After looking around on the web for what was available and what would be the best for us, I ended up deciding on Qdrant. A couple of other technical people recommended it as well, so it was decided that Qdrant would be the one..”

Mir Shahzad Mubeen

CTO at Honeycomb AI

[Read full review](#) 

“My main use case for Qdrant is as the vector database for an adaptive RAG assistant built with LangGraph and Azure OpenAI.

“Qdrant fits into my adaptive RAG assistant by serving as the core vector search engine within our LangGraph stateful workflow, storing embedded chunks of internal technical specifications, operational instructions, and regulatory documents alongside metadata payloads such as document form types. Text chunks are vectorized using Azure OpenAI's text-embedding-3-large model and indexed in Qdrant. During execution, a LangGraph retrieval node routes user queries to Qdrant to perform filtered vector search and reciprocal rank fusion across multi-query variations. Qdrant then quickly returns the top matching document payloads back to the LangGraph pipeline, which uses Azure OpenAI to evaluate context grounding and synthesize the final response..”

Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees

[Read full review](#) 

“I have been designing vector databases, so I use Qdrant mostly to explore how the engineers at Qdrant have built the whole vector database, specifically the multi-vector system. In terms of usage, I have been building a code intelligence tool. For example, I use embedding models from Cohere and store all the vector embeddings directly inside Qdrant. These are basic use cases without anything too fancy. I use Qdrant as a study case to help build other companies' internal vector databases.

“For extremely large code bases, I want to plot code and have the code generate summaries instead. I wanted to create a better mapping between memory.md files, vectorize them, and create a semantic relationship between symbols. By symbols, I mean function or variable symbols that I have extracted from tree-sitter and LSPs. I used those to create a much better semantic relationship between these things and provide them as context for Claude Code, which made my Claude Code more intelligent and also reduced token usage. These are basically what I call pre-processing or pre-context engineering. That is one use case where I have been using Qdrant and embedding models together along with my local coding stack..”

Adib M

[Read full review](#) 

Founder at a tech vendor with 11-50 employees

Setup

The setup process involves configuring and preparing the product or service for use, which may include tasks such as installation, account creation, initial configuration, and troubleshooting any issues that may arise. Below you can find real user quotes about the setup process.

“Technically, we are not really using Qdrant now, though we used it extensively for a long time, gaining better perspectives on vector stores and RAG systems for our own company, functioning as a help bot. We can query anything from Slack and retrieve data back for insights. When we were using Qdrant, we deployed it through Docker in a private cloud..”

Amrit Dash

[Read full review](#) 

Automation Engineer at a educational organization with 11-50 employees

“There is definitely a return on investment with Qdrant. Although I am on a free tier and therefore have zero investment, the results obtained have been excellent for a vector database. The time saved is substantial, with nearly three weeks or more for projects deployed with Qdrant Cloud in no-code platforms. The accuracy achieved through HNSW searching compared to others is notable..”

Chirag Morajkar

[Read full review](#) 

Lead Ai Tech And Tech Automation Engineer at a individual & family service with 11-50 employees

“My advice for others looking into using Qdrant is to start by prototyping for free using local Docker containers or Qdrant Cloud's free tier, as switching from local to production cloud requires only changing your API key and endpoint. Be sure to design your metadata payload schemas upfront and utilize payload indexing early so Qdrant can filter vectors during search without sacrificing speed. If you are using high-dimensional embeddings at scale, turn on Qdrant's native quantization to keep memory usage and cluster costs under control..”

Pawel Cislo

[Read full review](#) 

MLOps Engineer at a tech services company with 501-1,000 employees

Customer Service and Support

“I have not had any need to contact Qdrant's customer support. The documentation provided by Qdrant covers most queries effectively, and I foresee no need for support in the near future..”

Chirag Morajkar

Lead Ai Tech And Tech Automation Engineer at a individual & family service with 11-50 employees

[Read full review](#) 

“I rate the technical support of Qdrant as a nine because I think we have never reached out to them directly, but Qdrant has good support available online, and I can get answers from forums. The support is good..”

China Venkanna Varma Ponnamanda

Co Founder & CEO at SaYukth Private Limited

[Read full review](#) 

“I have not interacted much with customer support because I have not encountered significant issues from a developer's perspective that required me to reach out for help. My operations team also has not had many questions regarding pricing or usage, so I have not experienced major problems needing customer support..”

Validated user

Product Engineer at a tech vendor with 11-50 employees

[Read full review](#) 

“Qdrant's customer support is responsive and developer-focused. For free tier and open-source users, support is driven primarily through an active Discord community where core Qdrant engineers regularly step in to answer technical implementation questions. For cloud paying customers, support is handled via a dedicated ticketing system, which has been excellent in my experience..”

Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees

[Read full review](#) 

“I would rate Qdrant's technical support as community-driven. There's community support since we're not paying anything, and it's more the community support for it. It's open source, so we house it on our server.

“I think they provide enough information on the internet, and I am satisfied with it. They explain it well..”

AllanTan1

Chief Ai Scientist at Predictive Systems

[Read full review](#) 

Other Advice

“My advice for others looking into using Qdrant is to test it first by using Qdrant Cloud platform, which will give a very good idea of how it works and how it behaves. This way, if proceeding to a full enterprise version or another deployment option, it would be easier to work with. I have given this review a rating of nine out of ten..”

Mir Shahzad Mubeen

CTO at Honeycomb AI

[Read full review](#) 

“Setting up multi-query fusion still requires extra boilerplate code, payload metadata schema updates require custom migration scripts, and real-time visibility into HNSW graph indexing progress during bulk ingest could be better in the web UI. If you are using high-dimensional embeddings at scale, turn on Qdrant's native quantization to keep memory usage and cluster costs under control. I have rated Qdrant 9 out of 10 overall..”

Pawel Cislo

MLOps Engineer at a tech services company with 501-1,000 employees

[Read full review](#) 

“On a scale of one to ten, I would rate Qdrant overall as a ten. Regarding Qdrant's AI capabilities, thank you for asking about its governance and security. In my organization, Qdrant is deployed in the local server for my work, and it has really been great because the high performance has been very good for our work and for our use cases. The advice I would give to others looking into using Qdrant is that it is very good and they can always rely on it for high performance, especially for local AI development and implementation. I recommend it a lot..”

Femi Adeyemi

Senior Team Lead Systems Administrator at a comms service provider with 51-200 employees

[Read full review](#) 

“Qdrant is one of the best vector databases out there that is also open source. It is pretty good and has all the batteries I need. I think it is a great product for anyone looking to build their whole stack and also self-host them. Qdrant gives the flexibility to do it.

“Most vector databases support pretty much everything that a batteries-included full-fledged vector database should offer. However, the best thing about Qdrant is, first of all, it is open source. Second, it has very standardized client libraries for pretty much anything out there. It also benefits from being open source. Being open source and being pretty fast makes Qdrant stand out. There are plenty of open source databases out there, but Qdrant is the most complete one that is also pretty fast. These are the two things that make Qdrant a much better choice than any other vector databases.

“I mainly use Python and Rust, and Qdrant itself is written in Rust, which provides really good uses. I have used something like Chroma and Chroma is quite good, but it does not really scale for huge documents where Qdrant actually scales really well. That is why I call this the most complete package out there. Chroma is pretty good, but it does not really scale well. [Pinecone](#) is not open source, so it is also a bit of a hassle to self-host. Qdrant shines really well here.

“If someone needs something that is super fast for a tiny application, Qdrant might not be the best option. However, if someone is looking for scale or dealing with huge documents and medium to huge documents, Qdrant is a really good choice, especially if they want something that can be self-hosted and is open source. Qdrant is the best one out there for those needs. For smaller applications, someone should be looking into something Chroma or even DiskANN, which is by Microsoft. Qdrant is for medium to larger scale applications, not for tiny scale. For tiny scale, it is a bit of an overkill. My overall rating for Qdrant is seven point five out of ten..”

Adib M

Founder at a tech vendor with 11-50 employees

[Read full review](#) 

“Typically, we do not use proprietary databases such as [Microsoft Azure](#) or Horizon DB, and we have been using primarily open source databases. If you are talking about databases, primarily it is [PostgreSQL](#) or [MySQL](#). If you call us regarding vector databases, mostly it is Qdrant. Qdrant and Weaviate are partly paid and partly proprietary.

“I am a hands-on data scientist-cum-engineer, so I have used it myself.

“If you need to evaluate vector databases, Qdrant is among the very few databases where you can do a hybrid search. Qdrant is one of the few vector databases which does a hybrid search. By hybrid search, I mean searching through semantics, that is meaning, and also through word matching.

“Improving the model involves a lot of things, including hybrid search, reranking of results, and many other things are involved. Qdrant is a good and scalable vector database, and it is free. They have a free version. Their cloud services are paid, while the on-premise ones are free.

“Multiple query languages we have not used, primarily it is English. Multiple query languages would involve a different thing, such as translating the query and then searching the database, so that is a different ball game itself.

“In indexing and retrieving vectors, there are certain features which are not available as far as [Chroma](#) is concerned. You can store a feature of having metadata where, in vector databases or in RAG systems, things become more and more complex. In Qdrant, you can do that. I would rate this product an overall 8 out of 10..”

Uzair Faruqi

Head Of Data Science at Mjunction Services

[Read full review](#) 

“Regarding Qdrant's governance and security, I think it does a decent job because I have not encountered many problems with it. Qdrant runs as an HTTP plus gRPC

service, and my advice is never to expose it directly to the internet. Instead, put it behind an [NGINX](#) or API Gateway or use a VPN or VPC only access to protect your data from a network security perspective. On authentication or access control, since Qdrant is self-hosted, it lacks built-in role-based access control by default, which is an important limitation. Therefore, you must implement a reverse proxy, JSON Web [Token](#), OAuth 2, or API keys for network-level restrictions to enhance security.

“In terms of accuracy and reliability of output, Qdrant retrieves relevant vectors from the documents and chunks I provide, but it does not produce intelligent output, so it does not exhibit AI accuracy. Its accuracy depends on embedding quality, indexing quality, and query formulation quality. On reliability, I think it is very stable and consistent with high availability design, making its behavior predictable. Unlike LLMs, Qdrant does not hallucinate or guess answers; it only returns stored vectors. With the embeddings I have provided, it performs well and exceeds my expectations.

“One of the key limitations is that Qdrant does not have built-in role-based access control, and while being self-hosted is a benefit, it can also be improved. There is also no automatic governance policy, so I must design my own access rules, retention rules, or versioning. Additionally, there is no native encryption of payload fields to manage at the infrastructure or app level.

“If someone is looking for a vector database that can be easily hosted and integrated with Python to provide outputs quickly while being cost-effective, Qdrant would be a great option. It is also effective when handling large sets of data and fits well with LLM models such as LangChain, making it a top choice for anyone needing a vector database.

“Qdrant is an excellent vector database that anyone would want to use with RAG AI. I believe users will benefit significantly from Qdrant, and while they may explore other vector databases, I think they will notice a substantial difference with Qdrant compared to the alternatives. I would rate this product a 9 out of 10..”


Validated user

Product Engineer at a tech vendor with 11-50 employees

[Read full review](#) 


Top Industries

by visitors reading reviews




Company Size

by reviewers



by visitors reading reviews



Large Enterprise Midsized Enterprise Small Business

About this buyer's guide

Thanks for downloading this PeerSpot report.

The summaries, overviews and recaps in this report are all based on real user feedback and reviews collected by PeerSpot's team. Every reviewer on PeerSpot has been authenticated with our triple authentication process. This is done to ensure that every review provided is an unbiased review from a real user.

Get a custom version of this report... Personalized for you!

Please note that this is a generic report based on reviews and opinions from the collective PeerSpot community. We offer a [customized report](#) of solutions recommended for you based on:

- Your industry
- Company size
- Which solutions you're already considering

The customized report will include recommendations for you based on what other people like you are using and researching.

Answer a few questions in our short wizard to get your customized report.

[Get your personalized report here](#)

About PeerSpot

PeerSpot is the leading review site for cloud, AI, and business software. We created PeerSpot to provide a trusted platform to share information about software, applications, and services. Since 2012, over 22 million people have used PeerSpot to choose the right software for their business.

PeerSpot helps tech professionals by providing:

- A list of products recommended by real users
- In-depth reviews, including pros and cons
- Specific information to help you choose the best vendor for your needs

Use PeerSpot to:

- Read and post reviews of products
- Access over 30,000 buyer's guides and comparison reports
- Request or share information about functionality, quality, and pricing

Join PeerSpot to connect with peers to help you:

- Get immediate answers to questions
- Validate vendor claims
- Exchange tips for getting the best deals with vendor

Visit PeerSpot: www.peerspot.com

PeerSpot

244 5th Avenue, Suite R-230 • New York, NY 10001

reports@peerspot.com

+1 646.328.1944