


aws marketplace

Supabase

Reviews, tips, and advice from real users



Powered by  PeerSpot

Contents

Product Recap..... 3 - 4

Valuable Features..... 5 - 13

Other Solutions Considered..... 14 - 16

ROI..... 17 - 19

Use Case..... 20 - 26

Setup..... 27 - 29

Customer Service and Support..... 30 - 32

Other Advice..... 33 - 36

Trends..... 37 - 38

About PeerSpot..... 39 - 40

Product Recap



Supabase

Supabase Recap

Supabase Vector offers an efficient way to manage and query vector embeddings, catering to the needs of developers and data scientists seeking scalable solutions for vector-based data handling.

Supabase Vector is designed to streamline the process of storing, managing, and querying vector embeddings, essential for applications like machine learning algorithms and personalized recommendations. Its intuitive API and integration capabilities make it a preferred choice for tech professionals seeking a reliable backend for their vector data requirements. With flexible storage options and robust querying features, it accommodates the dynamic demands of AI-driven projects.

What are its key features?

- **Seamless Integration:** Connects effortlessly with databases, making it simple to incorporate into existing projects.
- **Scalable Storage:** Allows for efficient storage of large volumes of vector data, accommodating growth in usage.
- **Flexible Queries:** Supports complex queries, facilitating advanced data analysis and insights.
- **Real-Time Updates:** Ensures that data remains current with instantaneous updates.
- **Developer-Friendly:** Comes with comprehensive documentation and an intuitive API for ease of use.

What benefits or ROI should users look for?

- **Increased Efficiency:** Streamlines data management, reducing time spent on backend development.
- **Cost Effectiveness:** Offers scalable options that can save on infrastructure and maintenance costs.
- **Enhanced Accuracy:** Delivers precise query results, improving decision-making processes.
- **Improved Developer Productivity:** Simplifies complex tasks, allowing developers to focus on more critical activities.

Supabase Vector can be particularly beneficial in industries such as e-commerce for personalized product recommendations, in finance for fraud detection through pattern analysis, and in healthcare for patient data insights. Its capability to handle diverse sets of embeddings makes it versatile across different sectors needing robust data processing tools.

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- ✓ “Previously, 50% of searches returned no results or irrelevant products, but we changed that to 90% of searches returning relevant results, which led to a 30% increase in search-to-purchase conversions for the e-commerce client and was a big win for us.”



Boya Uday Kumar

AI Solutions Lead at ADP

- ✓ “Using Supabase Vector, I was able to set it up in a couple of days, got it validated, and started working on building up the algorithm on top of that.”



Pranay Kammari

Data Engineer at Caterpillar Inc.

- ✓ “Supabase Vector positively impacts my organization through major improvements in search relevance, faster development speed, a simpler architecture, and reduced engineering overhead, with better performance in real queries due to the hybrid SQL plus vector search.”



Verified user

Product Engineer at a tech vendor with 11-50 employees



“When I got to know about Supabase Vector, I fell in love with it since it worked for what I was using it for, and I just stuck to it.”



FavourEmosivbe

Automation Specialist at a consultancy with 11-50 employees



“The most valuable part using Supabase Vector is how complex things I was able to build and deploy in a short period of time.”



Alberto Hidalgo

Co-Founder at a tech services company with 1-10 employees



“We completed what we initially estimated as a fifteen-day job in just three days, thanks to the Supabase dashboards and triggers.”



Samir Patil

Co-Founder & CTO at Mango Giraffe



“Supabase Vector has positively impacted our organization as it is very convenient since our business databases are already hosted in Supabase, making integration easy.”



Daniel Pan

Director at a tech services company with 1-10 employees

What users had to say about valuable features:

“Supabase Vector offers a normal vector database solution like Pinecone and Weaviate. One of the things I particularly appreciate is how easy it is to add pgvector to your existing database by just running one line of SQL code in the SQL interface of your Supabase database. By default, you just run one line of code and then you have a vector database attached to your database, and everything is just set up. You do not have to do too much setting because you are already using the PostgreSQL instance for relational database.

“The ease of setting it up is primarily what stands out to me. In many other ways, I think it is just a state-of-the-art solution in the industry. Supabase Vector has made my work a lot easier, especially since we already use PostgreSQL instances on Supabase for our relational database..”

Victor Kola Omolaiye

AI/ML Engineer at a tech services company with 11-50 employees

[Read full review](#) 

“I appreciate Supabase Vector because it has RLS that allows its users to have complete access to just their data. With Supabase Vector, you have access to your data because of the RLS feature. You do not get to have access to other features and the database has structured tables, structured rows and columns. If I were building an app from scratch, I would have to build this back-end from scratch. However, with Supabase Vector, I do not have to do that. It serves that purpose. All I need to do is attach it with the front end of my website.

When I use it for RAG agents, instead of searching for a particular word for my RAG agent, it searches for related meanings. If I ask how to optimize customer support or how to make customer support faster, Supabase Vector will search for related meanings concerning these questions. It will give me answers relating to search bots because that would be related to customer support and helping to optimize the customer support experience. It will also give me related meanings that have to do with customer support. It is not searching for that particular keyword, but rather searching for all the related meanings. Because of this experience, it is much easier to use Supabase Vector for my work..”

FavourEmosivbe

Automation Specialist at a consultancy with 11-50 employees

[Read full review](#) 

“The best feature of Supabase Vector is definitely the Postgres. I already know SQL, so querying vectors with cosine simplicity and Euclidean distance feels natural. No new query language or learning is required. Even the setup is really easy; you just go to the dashboard, add a vector column there, and you're good to go. That's the best feature I've felt so far for Supabase Vector. The cost-efficiency, because everything lives in one managed database, helps with the cost as well. The performance has been solid so far, and the Supabase Vector client library makes installing, embedding, and searching really easy. I also appreciate the real-time features and the authentication specificity of Supabase Vector, which I feel is solid and secure.

“I initially thought it might take at least two to three months to set up a pipeline and start working. However, using Supabase Vector, I was able to set it up in a couple of days, got it validated, and started working on building up the algorithm on top of that. I am building a Python API suite on top of Postgres, so I am already refining my algorithms and searches on my databases to provide better results. I have become much more efficient than I initially estimated with Supabase Vector, and it is helping immensely..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“Supabase Vector offers many good features, beginning with semantic search. Another excellent feature is the integration of SQL with vector search, making it great for RAG, scalable indexing, and a simpler architecture. These are the main best features I find in Supabase Vector.

The SQL plus vector search combination has helped me significantly because it is easy to set up with PostgreSQL as it removes the need to choose between structured filtering and semantic search. By using both together in one query pipeline, vector search is smart, but not precise enough alone. Therefore, combining SQL for filtering first, and then applying vector search on the filtered set enhances the relevance, eliminating AI hallucinations and improving user intent handling.

Additionally, a key feature worth mentioning is recommendation systems. For instance, if a user says, 'Sweet vanilla perfume for winter,' SQL filters that database for winter and unisex requests, while the vector search finds vanilla-based scents with warm, sweet profiles, combining both signals effectively.

Supabase Vector positively impacts my organization through major improvements in search relevance, faster development speed, a simpler architecture, and reduced engineering overhead, with better performance in real queries due to the hybrid SQL plus vector search. This results in faster response times, lower database load, and more predictable queries, which together create practical metric improvements..”

Verified user

Product Engineer at a tech vendor with 11-50 employees

[Read full review](#) 

“The best features Supabase Vector offers include the ability to create a backend function or an API call, as I can use Supabase Vector's Edge Functions to implement those backend functionalities. Additionally, if I want to connect any OAuth provider like Google or LinkedIn, I can do that through Supabase Vector. It provides a seamless integration for that and it is really good if I want to integrate APIs, so it is a good vector database where I don't need to do much effort. Supabase Vector will do a lot of the running for me.

“Out of the features I mentioned, I find myself using Edge Functions in Supabase Vector the most because I don't need a dedicated backend like Node.js. In this case, I can use the edge function and create an API call. I also find it very useful if I want to establish security in my project, which is very important; Supabase Vector does a lot of the running. It also flags any security vulnerabilities, so I can provide and harden my security using RLS or other similar features.

“Supabase Vector has impacted my organization positively as it helps in strengthening the project without having dedicated backends, and it helps me establish a better workflow.

“Regarding the positive impact, I believe Supabase Vector helps in reducing costs and fastening the process because a lot of it can be done easily. It is easy to establish the database as there is a clear step-by-step process on how I can connect my project to Supabase Vector. So it's really helpful and seamless that way..”

Manrosingh Singh

Associate Developer at Deloitte

[Read full review](#) 

“Based on my experience with Supabase Vector, I personally feel that Unified Data Storage, where we store vector embeddings right next to our relational data in PostgreSQL, is one of the top features I appreciate. Additionally, there is another feature in Supabase Vector, which is SQL and Vector Hybrid Search, allowing us to combine traditional SQL queries with semantic search effortlessly. This is absolutely amazing because it lets us filter by metadata while doing similarity search. Another feature that stands out is the fast similarity search with multiple index types. Supabase Vector supports IVFFlat and HNSW indexes for fast approximate nearby neighbor search, giving us, as an agency, the ability to reduce response time to under 100 milliseconds, even with over 25,000 articles, which is a killer feature. The built-in authentication and security, with all the RLS policies applied automatically to vector queries, gives us trust with all the clients we work with. Supabase Vector's seamless integration with all AI tools, such as OpenAI, Hugging Face, LangChain, and Amazon SageMaker, is a significant edge for us. The free tier of 500 MB database with vector support included is a real help that Supabase Vector provides us. Additionally, we appreciate the auto-generated REST API, as we do not need to build any custom API for vector queries, making our development much faster. These are the features that are currently on my mind.

Supabase Vector has had a significant positive impact on the agency I work for by increasing client delivery. We built semantic search and AI chatbot features for over eight client projects in the past year, and these features are now standard offerings I provide to new clients as well. Revenue has increased by about 20% to 30% in value, and we reduced development time for semantic search projects from seven to ten days down to just two to three days because the auto-generated API and unified storage save us a lot of hours on back-end work. I saved over 20 hours of development time for one client project using Supabase Vector, which is a significant gain. Additionally, we have greatly improved client results, such as a 30% increase in search-to-purchase conversions for the e-commerce client I mentioned, which was a big win. There was also a 40% reduction in customer support questions after building the AI chatbot for the content management client I discussed, allowing the client to save about \$1,500 to \$2,000 per month on support staff. We have seen substantial improvements in search accuracy and scaled from handling over 10,000 articles to more than 25,000 articles, making

scalability another significant win. These are all the impacts that we have experienced in the organization, thanks to Supabase Vector. .”

Boya Uday Kumar

AI Solutions Lead at ADP

[Read full review](#) 

Other Solutions Considered

“I mentioned earlier that previously I used Pinecone. The primary reason why I have not completely switched, but why I started using pgvector is especially in cases where I already have a relational SQL database, PostgreSQL database for that project on Supabase. It just makes sense for me to use the vector database also..”

Victor Kola Omolaiye

AI/ML Engineer at a tech services company with 11-50 employees

[Read full review](#) 

“I evaluated other options, primarily Pinecone. The setup felt very challenging. I also considered Weaviate, which also had a heavy setup. By the time I was going through the setup information for Pinecone or Weaviate, I managed to set up Supabase Vector in a much faster timeframe and start running my projects..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“If you want to use Zilliz for instance, there is a whole lot of things that you have to do, a whole lot of database, a whole lot of code, a whole lot of APIs to call for different kinds of things. If you want to set up Pinecone, you have to go to the UI on the Pinecone website and pick the model, pick the chunk size, the dimension of your database, and a whole lot of other things. Pinecone is better than some others, but Supabase Vector is just straightforward..”

Victor Kola Omolaiye

AI/ML Engineer at a tech services company with 11-50 employees

[Read full review](#) 

“I did evaluate other options such as Pinecone, Weaviate, and Qdrant before deciding on Supabase Vector. I ultimately switched because of the comprehensive features it offers..”

Boya Uday Kumar

AI Solutions Lead at ADP

[Read full review](#) 

“Previously, I ran Postgres on top of AWS RDS, but that setup became expensive. Establishing my RAG-based vector system on AWS proved quite challenging and costly, particularly trying to connect Bedrock to Postgres within AWS. This led me to switch to Supabase Vector, which was much cheaper in that context..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“I did use a couple of other solutions before Supabase Vector, including manual data extraction with spreadsheets and separate data pulls from various APIs. I used Pinecone for one project and MongoDB with Atlas search for another. I switched to Supabase Vector because the manual processes were time-consuming, lacked real vector search capabilities, and the accuracy of results was poor. Transitioning to Supabase Vector enhanced search accuracy while providing unified storage without sync issues, proving to be more cost-efficient..”

Boya Uday Kumar

AI Solutions Lead at ADP

[Read full review](#) 

ROI

Real user quotes about their ROI:

“Currently, I am managing with fewer data engineers by using Supabase Vector. The time savings are notable; I expected to take two months to set up, but I accomplished it in just two days. That's over 50 days saved. Although I haven't seen major returns on investment yet, I look forward to realizing more profits down the line..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“I have not seen a return on investment yet because I am not investigating those points. However, I can say that building these kinds of applications with so many complex components in less than a couple of months is amazing. The use of these technologies definitely impacts reducing the time and cost of implementation or deployment. It has a really good impact on companies, employees, and finances..”

Alberto Hidalgo

Co-Founder at a tech services company with 1-10 employees

[Read full review](#) 

“I have seen a clear return on investment; we saved money because RDS is pretty costly. Supabase includes all functions we were using, and we previously took backups manually, which was acceptable. The dashboard's management made access straightforward for users and super easy to maintain, resulting in very few errors. We completed what we initially estimated as a fifteen-day job in just three days, thanks to the Supabase dashboards and triggers..”

Samir Patil

Co-Founder & CTO at Mango Giraffe

[Read full review](#) 

“The organization would not have noticed the impact because we work with a lot of persons that are not technical people, and even among the technical staff that we have at our organization, there are a lot of persons that are essentially full-stack developers for software who do not really use vector database the way I have to use it as an AI Engineer. However, for me, Supabase Vector has made my work a lot easier, especially since we already use PostgreSQL instances on Supabase for our relational database..”

Victor Kola Omolaiye

AI/ML Engineer at a tech services company with 11-50 employees

[Read full review](#) 

“I have seen a return on investment from using Supabase Vector, particularly through reduced development efforts, improved search efficiency, and better user engagement outcomes. The impact includes faster feature delivery and lower infrastructure costs.

We notice significant improvements when measuring metrics such as average response times, which have shifted from 800 milliseconds to 2.5 seconds down to around 200 milliseconds to 800 milliseconds, with click-through rates for recommendations improving by 45 to 70%. Additionally, search refinement has reduced redundant searches by 20 to 35%, and recommendation accuracy has improved typically by 25 to 60% for top results..”

Verified user[Read full review](#) 

Product Engineer at a tech vendor with 11-50 employees

“I have definitely seen a significant return on investment from Supabase Vector, saving about \$1,000 to \$1,500 per client project. Across eight client projects, I saved approximately \$4,000 to \$8,000 a month, and for the AI chatbot, I saved \$2,000 by eliminating support staff hours. Development time has drastically reduced, as the automatically generated REST APIs minimize the major development work. Instead of handling projects from seven to ten days, we now complete them in just two to three days. For the eight projects I have managed, we save about 160 to 240 total hours. There has also been a 30% increase in search-to-purchase conversions for our clients, along with 90% to 95% accuracy with the similarity search. Overall, development has accelerated, and the client value we deliver improves, leading to increased profits for our organization..”

Boya Uday Kumar[Read full review](#) 

AI Solutions Lead at ADP

Use Case

“Supabase Vector serves as the retrieval layer in my RAG application, where we store embeddings of our documentation. When a user asks, 'How do I reset my password?' we generate an embedding for the query and perform a similarity search in Supabase Vector to retrieve the most relevant documentation, passing that context to the LLM for an accurate response. This is primarily for an AI knowledge base, and while there are different use cases, this is the most common one.

In my workflow, Supabase Vector acts as the semantic retrieval layer. When new documents or product descriptions are added, I generate embeddings using an embedding model and store them in Supabase Vector along with the original data. When a user submits a natural language query, I generate an embedding for that query, perform a similarity search in Supabase Vector to retrieve the most relevant records, and then pass those results to the particular LLM, ensuring the model answers based on our own data rather than relying solely on its pre-trained knowledge. This workflow is primarily related to the RAG application..”

Verified user[Read full review](#) 

Product Engineer at a tech vendor with 11-50 employees

“At the initial stage, I used Supabase Vector in building AI chatbots and RAG agents. I used it as the vector store where I input knowledge bases and attach it to the RAG agent that acts as customer support for businesses. At present, I have started using it for web-coded apps. I basically use it as a back-end service where I can have a vector store, a back-end that is attached to the front end of the app and then I have something that serves as an authentication space for the app that provides authentication and authorization for the app users.

Supabase Vector works fine because for one of the apps I have used it for, it has had over 200 users and it actually worked fine. I have not had any issue to transfer to any other service at the moment. I have never thought of scaling or transferring to something else for a back-end because it serves the purpose of what I have used it for. I have basically used it for small teams, not large teams that need to cover thousands of users. Since it serves the purpose, I have never thought of scaling to something much more that will cover a wider range..”

FavourEmosivbe

Automation Specialist at a consultancy with 11-50 employees

[Read full review](#) 

“My main use case for Supabase Vector is for the projects that I am building, as I use Supabase Vector to host the databases, and whenever I'm integrating AI into my application as well. Supabase Vector provides a vector database which is good for LLMs, and it provides good embedding as well.

“One specific example of how I've used Supabase Vector in one of my projects is when I was working on a project where I had to build an RBAC matrix with different user roles, privileged user roles, and normal user roles. I used Supabase Vector to build these different tables and give different kinds of constraints on these tables. This ensures that only a normal user will be able to access the things that are applicable to them and also apply row-level security, which hardens the security of the system.

“When I have to implement an LLM into my project or an AI chat base needs to be implemented, I use Supabase Vector for the databases as it is a good vector database to apply the natural embedding of the data..”

Manrosingh Singh

Associate Developer at Deloitte

[Read full review](#) 

“As an AI Engineer, my primary use case of Supabase Vector is for storing vector databases that I use at retrieval and inference in my AI agent and RAG pipelines. I have been working with RAG solutions, and often times I have to chunk and embed information into vectors that need to be stored in databases so that at inference we can always retrieve the relevant chunks from the database.

“I was building on Pinecone, but then I had to test something with pgvector locally. After testing pgvector locally for a solution I built for Algorithmi UK, which is essentially a teacher co-pilot solution, I realized I had some information that I retrieve in raw text stored in buckets, some in SQL databases on Supabase, and then I had other ones that were semantic that I needed to run semantic search over. For those sections, I ended up testing locally with pgvector. When it was time to go live to production for internal testing or to staging for internal testing, I realized that instead of tearing down all of that pgvector coding and having to write the code for Pinecone, I could just as well use pgvector in internal testing.

“Recently, I started building a personalized portfolio website. I was not really impressed with all of the automated solutions out there, the frameworks that you have to put in specific information that grants a specific kind of output. I wanted to build a personal portfolio website where, beyond being able to pull my projects from GitHub, it is also able to have an AI-powered RAG solution on the back end that visitors can interact with and ask about me, ask if I have used a particular framework, if I have used a particular library, if I have built a solution that looks a certain way, and then it can answer without them having to contact me. That was my most recent use of Supabase Vector because the information that came with every README, every text file that I used for each project shown on my personal portfolio website had to be chunked so that the AI can use it in responding to user queries whenever they visit the website..”

Victor Kola Omolaiye

AI/ML Engineer at a tech services company with 11-50 employees

[Read full review](#) 

“I use Supabase Vector's Postgres for my projects, specifically vector-based projects for database usage. I also use it for my personal use on my personal apps that I build on the side. I work on an AI-based financial app, so I use Supabase Vector for that. The vector builds right into the top of that with `pg_vector`. The main use case that I use is building AI features, specifically the RAG pipelines, which are really augmented generation pipelines. I store text chunks and their embeddings for documentation and a knowledge base directly into the Postgres database. It lets me do search algorithms and regular relational queries, which is super convenient.

“In my most recent project, I've been using Supabase Vector for RAG-based pipeline generation. I store all my text chunks and their embeddings for documentation and knowledge as a knowledge base. I created a knowledge base in the Postgres database and use this RAG pipeline to run a search algorithm on top of the text-based database that I have created. So if a user gives a prompt, it will be easier for it to associate with the pre-existing text and figure out the best results for the user's search.

“As of now, that's the biggest thing I've been using Supabase Vector for lately.

“The pros of using the real-time features and authentication in Supabase Vector have been evident. It's been a quick setup compared to other databases such as Pinecone, which can be more cumbersome. With Pinecone, I have to manage a completely separate service, sync data in, and deal with extra costs and latencies. Supabase Vector keeps everything in one Postgres database, so I can join vector search results with my regular relational data table easily. It's a huge advantage for most apps.

“One thing I really appreciate about Supabase Vector is the backend capabilities. The Postgres changes over WebSockets allow you to subscribe to insert, update, and delete events on your table, including vector embeddings. For example, if a new document gets added to the embeddings generated, I can listen for those changes in real-time on the front end. The UI can instantly show updated search results or notify users when fresh content becomes available for semantic search. It works really well for collaborative apps such as dashboards or AI features where data is constantly updated..”

“I use Supabase Vector to build semantic search functionality for all the client websites that we handle, and it allows us to search by meaning rather than the exact keyword. As a team, we mainly use it for semantic search functionality, and our clients find it helpful too.

One of my clients had an e-commerce website for a fashion store, and customers were always frustrated because the traditional keyword search did not work really well. For example, if a user searched for summer clothes for the beach, they only got products that had exact keywords, but searching for light clothing for vacation would show completely different results. Many customers were frustrated about this problem and handed it over to us. We built a semantic search feature using Supabase Vector. I stored embeddings by creating vector embeddings for all 500 products using OpenAI's embedding API. Each product had the title, description, category, and tags, which we converted to vector format and stored in Supabase Vector. When a user types a query, the query is converted into a vector, and Supabase Vector finds the most similar products using similarity search and returns the results ranked by relevance. Previously, 50% of searches returned no results or irrelevant products, but we changed that to 90% of searches returning relevant results. This was a very significant improvement that our client did not expect from us, and it was only possible because we implemented that semantic search functionality on the website. Supabase Vector understood that walking shoes and hiking boots are related even though they are different words, and customers could now find whatever they needed in a more natural way, searching with natural language and getting the things they need. It was a great leap for that website, and it worked really well.

I have another project that I worked on where I worked with a client who had a large content management system with over 10,000 blog posts and articles. They wanted me to build an AI-powered chatbot that could answer user questions based on their content. I built a RAG and created vector embeddings for all 10,000 articles and stored them in Supabase Vector. When the user asks a question such as how to optimize a website for SEO, the query is converted to a vector and Supabase Vector finds the most relevant articles, mostly the top five. An AI model like an LLM reads those articles and generates a helpful answer that is sent back to the user. The

answer accuracy was about 85%, and the users were satisfied. It saved a lot of the client's time, and the client specifically mentioned to me that they saved about \$2,000 per month on support staff hours. This was a significant win for our agency, and all we did was store all the vector embeddings efficiently, and the similarity search was faster, scaling really well. The integration with other Supabase Vector features like authentication and a database made the whole system easier to build. Both projects were great successes, and Supabase Vector has been essentially the root cause for these successes.

We supported a small tech company with a knowledge base of over 5,000 support articles where users needed to search for help articles relevant to their queries, filtering for articles from the last six months and from specific product categories such as mobile app or desktop app. This posed a major challenge for us. Without hybrid search, if a user searched for help on resetting a password in the mobile app, they would see all relevant articles across the entire database, potentially including outdated articles. We implemented a hybrid approach using SQL and semantic search. We wrote a query for the semantic search to find similar articles for resetting passwords in the mobile app while the SQL filters ensured that we only returned articles from the mobile app category created in the last six months. This approach successfully retrieved five relevant articles with an accuracy of almost 95%. Users found it very useful and received the right answers. Narrowing down specific details was only possible with this hybrid search. .”

Boya Uday Kumar

AI Solutions Lead at ADP

[Read full review](#) 

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.

“I initially thought it might take at least two to three months to set up a pipeline and start working. However, using Supabase Vector, I was able to set it up in a couple of days, got it validated, and started working on building up the algorithm on top of that..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“My experience with Supabase Vector's pricing, setup cost, and licensing is that it was pretty straightforward; it has an MIT license, which is good for us. We generally use a \$200 to \$300 subscription based on our database size, which is also fine for vectorization and other tasks..”

Samir Patil

Co-Founder & CTO at Mango Giraffe

[Read full review](#) 

“Supabase Vector is very easy to deploy. You do not have to go through a lot of stress to deploy it and you do not even need to have technical knowledge to deploy Supabase Vector. If you go on YouTube, and you are a newbie trying to deploy Supabase Vector, you can get on YouTube and since it is open source, it makes it even easier to deploy. You can go on YouTube, check out the videos or a tutor can guide you on how to deploy. It is easier. You do not need so much technical knowledge for you to actually deploy Supabase Vector..”

FavourEmosivbe

[Read full review](#) 

Automation Specialist at a consultancy with 11-50 employees

“The initial setup is straightforward compared to other products.

Deployment took less than ten minutes, and configuration and setup took about 15 to 20 minutes. It was even quicker for cloud services as it involved subscribing and commissioning the service.

I rate the process an eight and a half..”

KolawoleMobolaji

[Read full review](#) 

Managing Director at TEVC Concept

“Integrating Supabase Vector into our existing systems is not complex; it was relatively straightforward due to its direct build on PostgreSQL, fitting naturally into our workflows without needing a separate vector service setup.

The most challenging part of implementing Supabase Vector was designing an effective embedding strategy, as that presents difficulties common to all vector databases. Additionally, chunking and structuring data for heavy use cases posed challenges, particularly around embedding consistency during product data updates..”

Verified user[Read full review](#) 

Product Engineer at a tech vendor with 11-50 employees

“The integration and setup were complex, but the pricing is decent. I am using the Supabase free trial, which allows two databases, but it would be nice to have more than those two. The biggest complexity lies in the knowledge required to get all these pieces together. The app I have implemented uses several technologies including KMP for building Android and iOS apps, Firebase, Supabase, n8n, webhooks, and everything built on GitHub. You have to have decent knowledge of all those technologies and how to handle the Google API console to create the APKs for the embeddings. The knowledge requirement is the most complex part, but apart from that, the pricing is not a significant issue..”

Alberto Hidalgo[Read full review](#) 

Co-Founder at a tech services company with 1-10 employees

Customer Service and Support

“We usually just check the documentation or ask artificial intelligence for tech support, which is why we do not have much experience with customer support..”

Daniel Pan

Director at a tech services company with 1-10 employees

[Read full review](#) 

“Customer support is handled using emails at the moment, and since I am not using customer support from Supabase, I suppose it does not have any problem or concern in that regard..”

Alberto Hidalgo

Co-Founder at a tech services company with 1-10 employees

[Read full review](#) 

“The customer support for Supabase Vector has been good. I find their documentation very helpful, and I have not had to contact customer support directly since most information is available there. Even Stack Overflow has provided a wealth of data for addressing bugs and challenges I faced during development..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“The customer support has been excellent for my use as a new agency owner. Community support from helpful developers and engineers provides fast responses on GitHub issues and community forums. I have not encountered significant issues with customer support. I would rate the customer support a 9 out of 10..”

Boya Uday Kumar

AI Solutions Lead at ADP

[Read full review](#) 

“My experience with customer support is positive; I needed to reach out to Supabase, and the customer support team was pretty good, responding quickly to my issue. They resolved tickets efficiently, which I found very productive.

“I would rate the customer support a nine since they replied quickly and answered my questions properly, which helped me a lot..”

Samir Patil

Co-Founder & CTO at Mango Giraffe

[Read full review](#) 

“I have been able to figure out my issues on my own. When I have an issue, I tend to reach out to either a community of people where we do automations. If I have an issue, I could drop it in the community and somebody is always there to guide me. Somebody who may have gone through that challenge before is always there to help. I have always been able to solve it out with the help of my community or sometimes YouTube, but I have never had to contact support. Over time, I have gone through that stage where I even have to contact anybody because I am already used to Supabase Vector because I have used it for a long time..”

FavourEmosivbe

Automation Specialist at a consultancy with 11-50 employees

[Read full review](#) 

Other Advice

“For moderate usage, Supabase Vector is quite reasonable. I would suggest that others consider using Supabase Vector, as I believe it offers significant benefits. I would rate this product a 9 out of 10..”

Pranay Kammari

Data Engineer at Caterpillar Inc.

[Read full review](#) 

“My advice for others looking into using Supabase Vector is to take into account how you are going to create the embeddings, as it is not an option implemented straight away in Supabase. You will need to handle the embeddings with an external tool, so make sure to consider how these integrations are going to be made. I would rate this product an 8 out of 10..”

Alberto Hidalgo

Co-Founder at a tech services company with 1-10 employees

[Read full review](#) 

“My advice for others considering Supabase Vector is to start simple and avoid over-architecting early on. It is crucial to design your data model carefully and always use SQL filtering first before vector search to enhance results.

Overall, Supabase Vector is a good product that needs some improvements, but it has many redeeming qualities. It is easy to integrate, especially if you are already using Supabase for your database schema. I encourage others to try Supabase Vector as it proves to be a reliable tool. I would rate this product a 9 out of 10..”

Verified user

Product Engineer at a tech vendor with 11-50 employees

[Read full review](#) 

“With a PostgreSQL instance on Supabase to hold user data, hashed passwords, and all of those things, and then having to go elsewhere, say Pinecone or Weaviate, to go and create a vector database, to use all of those in my codebase means I have to connect to them independently. I would have a framework for connecting to the Pinecone vector database, have a framework for connecting to Supabase Vector and to the Supabase relational database, and all of that. But now I can just create everything on Supabase, copy my project URL, and just hook it up and everything just works. This reduces complexity and helps saves time.

“Documentation would be nice, but with the advent of AI-assisted coding, the accuracy of Supabase Vector is good. I would rate this review a 7..”

Victor Kola Omolaiye

AI/ML Engineer at a tech services company with 11-50 employees

[Read full review](#) 

“I would rate Supabase Vector an 8.5 overall. I would round that to a 9 because all

the strongest features such as unified storage, hybrid search, and fast similarity search that we usually use are killer features. Supabase Vector saves a lot of development time for our agency and is essential in terms of semantic search and RAG systems, which gives me about six points. The remaining three points come from the free tier supporting millions of embeddings at no cost, which is a significant benefit for all agencies. I decrease one point due to minor documentation gaps for advanced features and the absence of built-in embedding generation, leading me to choose a score of nine. My overall review rating for Supabase Vector is 9 out of 10.

My advice to others looking to use Supabase Vector is to start with a project, build things, and get to know the features within Supabase Vector, especially exploring semantic search. This approach provides ample opportunity to understand how Supabase Vector works while working with substantial amounts of data. Therefore, I encourage everyone to begin with a small project, manage a use case, and once it delivers returns, consider transitioning to the Pro plan if necessary. It is advisable to switch from Pinecone or [MongoDB Atlas](#) search to Supabase Vector for an extensive hands-on experience that significantly reduces development time. .”

Boya Uday Kumar

AI Solutions Lead at ADP

[Read full review](#) 

“On a scale of one to ten, I would rate Supabase Vector as eight out of ten.

“I chose eight out of ten because, as I said, it has a lot of capabilities and it is a very useful tool to use, so I would give a very strong rating of eight.

“Regarding Supabase Vector's AI capabilities, I feel the security that Supabase Vector provides is really helpful and it really strengthens its capabilities. If one is talking about its support center and AI features used there, then it is good as well.

“When it comes to Supabase Vector's AI capabilities, I find its output pretty

accurate since I've been using Supabase Vector to integrate AI in my project. I find it pretty reliable and it is easy to integrate seamlessly into the software development life cycle as well.

“I'm not sure how Supabase Vector was purchased; that is some other department's job, but this is how I use it.

“I think I have mentioned everything that I wanted to mention.

“My advice for others looking into using Supabase Vector is that the first learning curve might be a little challenging, but I would say spend as much time in the beginning as possible to learn all the things that Supabase Vector can do, how to connect your projects to Supabase Vector's API vectors, and learn everything in the beginning so you can seamlessly establish a process that works for you.

“I feel Supabase Vector is doing a very great job in reaching out to companies and helping them. I believe every new entrepreneur who wants to build their own product, their own software product, should be using Supabase Vector, and a lot of them do as well. I think it's the best vector database out there. My overall rating for Supabase Vector is eight out of ten..”

Manrosingh Singh

Associate Developer at Deloitte

[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