

aws marketplace

Cohere

# Reviews, tips, and advice from real users



Powered by  PeerSpot



# Contents

- Product Recap..... 3 - 4
- Valuable Features..... 5 - 10
- Other Solutions Considered..... 11 - 12
- ROI..... 13
- Use Case..... 14 - 16
- Setup..... 17 - 18
- Customer Service and Support..... 19
- Other Advice..... 20 - 22
- Trends..... 23 - 24
- About PeerSpot..... 25 - 26

# Product Recap



# Cohere Recap

Cohere provides a robust language AI platform designed for efficient implementation in various domains, offering advanced features for automation and data analysis.

Cohere delivers a scalable AI language model that facilitates automation in data-driven environments. Highly adaptable to industry-specific requirements, it supports tasks such as text generation, summarization, and anomaly detection. This flexibility, along with its integration capabilities, makes it valuable for tech-savvy users seeking seamless AI solutions.

## What are the notable features of Cohere?

- **Adaptive Language Models:** Efficiently handle diverse text-based tasks with customizable models.
- **Data Integration:** Seamless integration into existing databases and applications for smooth operation.
- **Scalability:** Support for scaling processes to meet enterprise-level demands without performance loss.

## What benefits should users consider in reviews?

- **Cost Efficiency:** Reduction in manual task overhead leading to budget savings.
- **Performance Improvement:** Boost in operational productivity due to automation and precise data handling.
- **Customization:** Ability to tailor AI functionalities to specific industry needs or unique processes.

Cohere sees significant use in finance, healthcare, and marketing, enabling precise data analysis and strategic insights. In finance, it assists with detecting market trends, while in healthcare, it supports clinical documentation and research analysis. Marketing uses include content creation and consumer sentiment analysis.

# Valuable Features

Excerpts from real customer reviews on PeerSpot:

- ✔ “Cohere's Embed English v3.0 is a cloud-hosted model that took less time to embed the textual data and was more than 50 to 60% faster than other models, even somewhat faster than text-embedding-3 from OpenAI, helping to reduce development and embedding times.”



**AmanSingh5**

Engineer at Roche

- ✔ “When it creates a new test, it creates it almost 70 to 80% correctly without errors; the time savings are significant—what previously took one or two days can now be completed in two to three hours maximum.”



**Gokul Anil**

Sr Test engineer at a tech vendor with 10,001+ employees

- ✔ “Cohere helped us with all three aspects: money is saved, time is saved, and we needed fewer resources to meet our end goals.”



**Shivam Singh**

Senior Solution Architect at Hitachi Systems India Private Ltd

- ✔ “A key advantage of integrating Cohere’s reranking model is that it aligns with client requests to include a reranking module — a widely recognized method for improving RAG quality. Additionally, the API demonstrates strong performance in terms of response speed.”



**Daniel Pan**

Founding Engineer at Agentize.AI

- ✔ “Cohere has positively impacted my organization by helping our customers work more efficiently when creating requests, and the embedding results are of very high quality.”



**Rustam Sharipov**

DevOps Engineer at CHI Software

- ✔ “Cohere has helped my organization innovate and stay ahead in our industry as Cohere was better than Titan, and it helped us to secure the client's confidence and we moved from proof of concept to production.”



**Diganth Sanghvi**

Ai engineer at a tech vendor with 10,001+ employees

- ✔ “I assess the value of Cohere's API support in my business operations as easy to integrate.”



**Verified user**

Manager at a insurance company with 51-200 employees

What users had to say about valuable features:

“The best feature Cohere offers is the Reranking model.

What stands out for me about the ranking model is that it improved performance in my work.

Cohere positively impacted my organization by improving the performance of my RAG system.

I noticed a 10% improvement in my log system after using Cohere..”

**Verified user**

Senior Data Scientist at a tech vendor with 10,001+ employees

[Read full review](#) 

---

“I assess the value of Cohere's API support in my business operations as easy to integrate.

The specific benefits I have seen from using Cohere include saving time to summarize information and for chat completion and related tasks.

I use Cohere for chat completion purposes.

My thoughts on the summarization feature are that it is beneficial for impacting our data analysis tasks..”

**Verified user**

Manager at a insurance company with 51-200 employees

[Read full review](#) 

“I believe Cohere offers excellent features, especially the cloud-hosted model and the API calls. The number of times I can call the API within a minute is very good. The ping is great; I have started a request to Cohere model, and it was very quick to respond. The best part was the free tier because most models do not provide a free tier.

Regarding benefits, Cohere is less costly than other models. If I talk about OpenAI or Google embedding models, they charge highly compared to Cohere. Regarding the training data, Cohere has the most data embedded or trained with the most English. Cohere's Embed English v3.0 has been trained with much more data than other models, including OpenAI. This gives an extra benefit to my organization..”

**AmanSingh5**

Engineer at Roche

[Read full review](#) 

---

“The best features Cohere offers are that it is fast and great.

“Speed has helped me in my day-to-day work, and I really notice the difference because it responds very quickly to LLM requests.

“Cohere has positively impacted my organization because I use it with Oracle, and in an enterprise way, it helped me offer clients a unique place to develop and use LLM. I can tell you that it helped me offer clients a unique place to develop and use LLM, as I use Oracle services..”

**Verified user**

Senior Data Scientist at a tech vendor with 10,001+ employees

[Read full review](#) 

---

“In my opinion, the best features Cohere offers are the embedding flexibility and the normal way the LLM reacted to the embeddings of Cohere. I used OpenSearch to integrate and store all the embeddings, and I used Titan as well to store the embeddings in OpenSearch, but the result was much better.

The flexibility I mentioned is evident because when we were using Titan, it was hallucinating a lot and not giving proper answers because I felt the embedding was poor. When we used it with Cohere, the embeddings were better and the chatbot with the LLM that used the embeddings from Cohere answered in a better way.

Cohere has positively impacted my organization as the project was a success. Clients were really happy with the results, and we received more business from them..”

**Digant Singhvi**

Ai engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“Cohere's command model is particularly useful for scenarios where consistent controlled output is more important, especially where we need creative responses, so I think Cohere's command model fits better in that case. We also found it well suited for structured enterprise tasks such as policy drafting, knowledge extraction, and generating standardized text for operational workflows.

It struck a good balance between fluency and predictability, which helps our team and is valuable for our business-critical applications, giving better insight to our team.

One of the major benefits I saw was data isolation and governance since Cohere has been implemented.

Consistent output quality, strong instruction following, and excellent embedding performance for retrieval tasks have benefited our organization. It was also offered from Amazon Bedrock, so this complete offering and strength from Cohere's command model helped our customers, and it is enterprise-friendly with deployment options such as VPC and data isolation that helped significantly.

Data privacy was a major concern because we operate from Asia-Pacific, and there is strong governance for data privacy in our country, so data privacy is the major compliance that helped us here..”

**Shivam Singh**

Senior Solution Architect at Hitachi Systems India Private Ltd

[Read full review](#) 

# Other Solutions Considered

“I did not previously use a different solution; this was the first time we were developing it and Cohere was from the start a competitive alternative to Titan, so we were always going to choose between Cohere or Titan..”

**Diganth Sanghvi**

Ai engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

---

“I have experience working with several different AI products, but I do not perceive any significant difference between them; they are nearly identical with accuracy varying slightly across certain areas..”

**Verified user**

Manager at a insurance company with 51-200 employees

[Read full review](#) 

---

“Before choosing Cohere, I evaluated AWS Titan.

The key factors that led me to choose Cohere over other AI development platforms are that I was experimenting, and Cohere was better than Titan..”

**Diganth Sanghvi**

Ai engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“Previously, I was using all-MiniLM-L6-v2 and switched to Cohere because all-MiniLM-L6-v2 needed to be locally deployed. That model was processing locally, and the results I was getting from that model, even though it was open source, I was not satisfied. That is why I switched to Cohere..”

**AmanSingh5**

Engineer at Roche

[Read full review](#) 

---

“I have evaluated OpenAI's Embed English v3 and text-embedding-3 models. I have evaluated multiple models, and I even evaluated some models from Hugging Face..”

**AmanSingh5**

Engineer at Roche

[Read full review](#) 

---

“We have tried multiple models, but we found that Cohere's command was a better fit for our needs.

We explored models from Anthropic and AWS native models such as AWS Titan Text before choosing Cohere..”

**Shivam Singh**

Senior Solution Architect at Hitachi Systems India Private Ltd

[Read full review](#) 

# ROI

Real user quotes about their ROI:

“I do not have relevant metrics about the return on investment from using Cohere yet because the customer's application is in a paging stage and has not been released. However, I understand that it is performing well, and we plan to continue with it..”

**Rustam Sharipov**

DevOps Engineer at CHI Software

[Read full review](#) 

---

“I can highlight two benefits. Cohere charges less than OpenAI, so it saves cost. In the second use case, the timing is significant. Cohere's Embed English model took less time to embed than OpenAI's embedding ada-002 model. In this case, it also saves time. These two benefits I can highlight..”

**AmanSingh5**

Engineer at Roche

[Read full review](#) 

# Use Case

“My main use case for Cohere is Retrieval Augmented Generation.

A specific example of how I use Retrieval Augmented Generation with Cohere is for information retrieval systems..”

**Verified user**

Senior Data Scientist at a tech vendor with 10,001+ employees

[Read full review](#) 

---

“I work with Cohere and have been doing so for about two months.

Currently, I am working with AWS Cloud and cloud services, and we use models like GPT-4o mini, 2.1, and Cohere.

We primarily use English only, with no other languages..”

**Verified user**

Manager at a insurance company with 51-200 employees

[Read full review](#) 

“My main use case for Cohere is for LLM and chatbot development.

“I use Cohere to fill boxes about documents, specifically about tenders.

“Cohere helps me fill boxes about documents, and I work with docx documents for a private company..”

**Verified user**

[Read full review](#) 

Senior Data Scientist at a tech vendor with 10,001+ employees

---

“My main use case for Cohere is that it's a good embedding model. I have used it with Titan, but Cohere came out better.

A specific example of how I've used Cohere for embeddings is when I was working with one of our clients where we were establishing a chatbot that can help us go through 31 PDFs. For embedding, we used Cohere and Titan, and Cohere was a superior product.

I have integrated Cohere in that chatbot project using SageMaker, and it was an easy API call that I used..”

**Digant Singhvi**

[Read full review](#) 

Ai engineer at a tech vendor with 10,001+ employees

---

“We adopted Cohere primarily for their command model to support enterprise-grade text generation and NLP workflows.

There was a use case for one of our customers where they required automated text generation and summarization of long documents and draft creation for internal content, so we used Cohere's command model with AWS Bedrock.

For another customer, there was a similar use case but they also wanted semantic search and RAG, and instruction-based responses for chat and workflow automation were required, so we used Cohere's command model for that..”

**Shivam Singh**

Senior Solution Architect at Hitachi Systems India Private Ltd

[Read full review](#) 

---

“I have used Cohere in a RAG use case where I had to vectorize some data. I used multiple models in RAG to find a better model that could give superior results. I was trying to find a cloud-hosted model, and Cohere's Embed English v3.0 is a cloud-hosted model that took less time to embed the textual data. When I was trying to get the similarity search after embedding that data, Cohere provided much better results.

Let's suppose I had to embed 100 documents at a time. Most other models, including all-MiniLM-L6-v2, took more time when I was trying to embed using that model. When I tried Cohere, it was much faster. I would say it was more than 50 to 60% faster than those models. It was even somewhat faster than text-embedding-3, which is from OpenAI. So Cohere helped to reduce the development time and embedding times..”

**AmanSingh5**

Engineer at Roche

[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'm more focused on the speed and overall quality of the model itself and the chat flow as a whole solution. That's why I'm not in the position to comment on the price and setup cost as there are DevOps working on this piece..”

**Daniel Pan**

Founding Engineer at Agentize.AI

[Read full review](#) 

---

“It's very easy. You just need an API key, and all the configurations are there. It's very easy to start.

If you have worked with something that requires API keys, you should be good to go. I don't think you need a lot of experience..”

**CollinsOmondi**

Mobile Application Developer at Uamuzi Foundation

[Read full review](#) 

---

“The setup was pretty smooth. I was able to find things easily. The documents were readily available on the internet, and I was able to find and integrate them without any issues. I subscribed to emails about new model updates, which allowed me to stay current. Oracle has now wrapped it inside their own AI, and we are using the latest version of Cohere as our chosen model..”

**Gokul Anil**

Sr Test engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

# Customer Service and Support

“We haven't had any issues to escalate to Cohere's support because reranking is an optional feature in our product, and we haven't seen any significant issues so far..”

**Daniel Pan**

Founding Engineer at Agentize.AI

[Read full review](#) 

---

“I have not interacted with Cohere's support team. However, I contacted Azure about the slowness, and we decided to use smaller chunks of information during the embedding process..”

**Rustam Sharipov**

DevOps Engineer at CHI Software

[Read full review](#) 

# Other Advice

“I intended to clarify that I use Cohere for chat completion.

My primary concern stems from the missing OCR capabilities.

My overall review rating for Cohere is seven out of ten..”

**Verified user**

Manager at a insurance company with 51-200 employees

[Read full review](#) 

---

“My advice to others looking into using Cohere is to try it.

My company does not have a business relationship with this vendor other than being a customer.

I gave this review a rating of 8..”

**Verified user**

Senior Data Scientist at a tech vendor with 10,001+ employees

[Read full review](#) 

“I do not want to add anything else about the features, including anything about accuracy or ease of use.

“I do not have specific advice to give to others looking into using Cohere. I gave this review a rating of 6..”

**Verified user**

Senior Data Scientist at a tech vendor with 10,001+ employees

[Read full review](#) 

---

“Cohere has helped my organization innovate and stay ahead in our industry as Cohere was better than Titan, and it helped us to secure the client's confidence and we moved from proof of concept to production.

The advice I would give to others looking into using Cohere is to go for it. Cohere is the best embedding model, and I would not recommend wasting time with AWS Titan.

I rate Cohere a 10 out of 10 because Cohere was way better compared to Titan..”

**Digant Singhvi**

Ai engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

---

“Cohere provides a free tier, and any developer who is starting their journey can use Cohere for RAG use cases. They can utilize the model benefits. After using Cohere, I got distances after the similarity search that were much lower compared to other vectorization and embedding models. The only model that performed better than Cohere was OpenAI's text-embedding-3-large. It was good, but Cohere was the second-best performing model in my use case.

I think Cohere's use cases are excellent, and I would suggest Cohere to others because of the less response time and time-saving in the process. It is also cheaper than other models. I would give this review a rating of eight out of ten..”

**AmanSingh5**

Engineer at Roche

[Read full review](#) 

---

“Data privacy was a major concern because we operate from Asia-Pacific, and there is strong governance for data privacy in our country, so data privacy is a major compliance that helped us here.

Cohere offers great customization options.

If governance, consistency, and data privacy are priorities, Cohere meets our organization's requirements well.

I recommend that anyone, especially in environments where governance, consistency, and data privacy are priorities, should choose Cohere, particularly the command model for teams looking for a controlled enterprise-safe alternative for text generation, summarization, and instruction automation.


Currently, we have used Cohere from the AWS Bedrock offering only, but since AWS has changed their third-party model availability from partner accounts, in the future, we are going to be a reseller for Cohere.

The documentation and learning resources were very helpful.

Our overall review rating for Cohere is 8 out of 10..”

**Shivam Singh**

Senior Solution Architect at Hitachi Systems India Private Ltd

[Read full review](#) 

# Top Industries

by visitors reading reviews

Financial Services Firm



Manufacturing Company



Comms Service Provider




Marketing Services Firm




# Company Size

by reviewers



by visitors reading reviews



Large Enterprise

Midsize 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 software running on AWS and other platforms. 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](http://www.peerspot.com)

## PeerSpot

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

[reports@peerspot.com](mailto:reports@peerspot.com)

+1 646.328.1944