

aws marketplace

**Cohere Command R**

# Reviews, tips, and advice from real users



Powered by  PeerSpot



# Contents

Product Recap.....	3 - 4
Valuable Features.....	5 - 9
Other Solutions Considered.....	10 - 11
Use Case.....	12 - 15
Setup.....	16
Customer Service and Support.....	17
Other Advice.....	18 - 19
Trends.....	20 - 21
About PeerSpot.....	22 - 23

# Product Recap



Cohere Command R

# Cohere Command R Recap

Cohere Command R offers a sophisticated suite of AI-driven tools tailored for advanced data analysis and language comprehension, making it a go-to option for businesses seeking efficiency.

Cohere Command R stands out as a versatile AI tool equipped to handle complex linguistic tasks and streamline data operations. It's designed to support developers and businesses in deriving actionable insights through advanced AI capabilities. With its robust performance, it caters to the growing needs of enterprises looking for cutting-edge technology in data processing and language-related tasks.

## What are the key features of Cohere Command R?

- **AI Language Models:** Provides state-of-the-art AI models for natural language understanding and generation.
- **Scalability:** Easily adapts to increasing data loads, ensuring sustained performance.
- **Integration Capabilities:** Seamlessly combines with existing IT frameworks for enhanced efficiency.

## What benefits should users consider in reviews?

- **Enhanced Efficiency:** Reduces manual workload, allowing focus on strategic tasks.
- **Cost Reduction:** Minimizes operational costs through automation and streamlined processes.
- **Improved Accuracy:** Delivers precise insights, boosting decision-making processes.

In industries like finance and healthcare, Cohere Command R is implemented to automate data-heavy processes and improve analytical capabilities, leading to increased productivity and accuracy. It supports professionals in achieving precise outcomes, enhancing overall industry effectiveness.

# Valuable Features

Excerpts from real customer reviews on PeerSpot:

- ✓ “After this model release, when we integrated this model on our platform, around 20% of users came to use chatbot, and previously they were facing complaints that the chatbot replied too slowly or hallucinated a lot, but after using this model the complaints are very minimal and their support tickets are reduced by 5% to 10%.”



**Husain Barwala**

AI Engineer at Walkover Web Solutions

- ✓ “Personally, compared to other models, Cohere Command R is pretty easy to set up and good for what I need as of now.”



**Collins-Omondi**

Mobile Application Developer at Uamuzi Foundation

- ✓ “After implementing Cohere Command R, the whole process became streamlined, reducing time and increasing end user engagement.”



**Singh Aman**

Generative AI Engineer at Tata Consultancy

- ✔ “The best feature Cohere Command R offers is the latency, which is faster than other solutions I have tried and has improved the latency and our time to delivery.”



**Verified user**

Engineer at a construction company with 11-50 employees

What users had to say about valuable features:

“Cohere Command R works for what I need. I know there are many other models and many other free models. I have tried CodeGemma, but it is not for what I was trying to do; it is more about coding. I wanted something interactive, focusing on the language side of it, not the coding side.

Personally, compared to other models, Cohere Command R is pretty easy to set up and good for what I need as of now.

Deploying this solution is pretty similar to working with any other model for me. I cannot really say much about complexities, but I am a bit technical, so the process is quite the same..”

**Collins-Omondi**

Mobile Application Developer at Uamuzi Foundation

[Read full review](#) 

“The most valuable feature is its strength in RAG. Cohere Command R performs well with prompts that include written context and expected outputs, mostly grounded in that context. The built-in citation capability is very useful as it makes answers easier to verify and increases user trust. I also value the low latency and throughput compared with some larger general-purpose models. Cohere Command R feels more practical for production-scale chatbots and documentation Q&A workflows where response time matters. It was very quick. The best part was the 128k context window, which is helpful when working with longer documentation and multiple chats.

“After implementing Cohere Command R, the whole process became streamlined, reducing time and increasing end user engagement. The increase in user trust was due to the reliable grounding of responses with citations. It has literally improved the quality of document-based AI responses in RAG workflows. Users are less likely to receive generic or unsupported answers because the model can ground responses with the retrieved answers and citations.

“Cohere Command R also helps accelerate prototyping. Instead of spending excessive time tuning prompts for basic Q&A, I can get a usable baseline quickly and then focus on retrieval quality, evaluations, and guardrails. For internal knowledge assistants, this reduced friction and improved trust in the generated outputs..”

**Singh Aman**

Generative AI Engineer at Tata Consultancy

[Read full review](#) 

“Cohere Command R is a very good model that I have used in past scenarios. It is a chat-focused language model, like a workhorse conversational AI model. I use it as a Retrieval-Augmented Generation tool because when I upload a large document to this model, it is able to answer me with a very accurate response compared to other models. It is good at providing grounded answers. You can give it some documents and ask a question and it answers using those documents. Crucially, it also tells you which document each piece of information came from, which is known as citations. I appreciate that very much and we have been using it in our product as well.

“My company has been using this model for Retrieval-Augmented Generation purposes for six months. This model is also fast compared to some other models when talking about RAG, Retrieval-Augmented Generation. The context window of this model is also good at around 128,000 tokens. I appreciate this very much.

“I can provide you with some pros of this model that we appreciated, which is why we integrated it. This model is very cheap with an input cost of 0.15 and an output cost of 0.6 per million tokens. The context window is very high so that large documents are able to fit with this model. Built-in citations was a very good feature that I appreciate, as RAG responses come with resource attributions from the document. The model identifies which part of the document was picked and according to that document part, the model is answering. I appreciate this very much and our users appreciate it as well. They were able to get confidence that this model is giving 100% accurate answers according to the data present in the document. Beyond this, this model speaks more than 10 languages, though our users are mostly using three to four languages. This model is fast. The structured output capability is also good as it returns reliable JSON. If you are working with a production level pipeline with documents, this model reads the document very carefully and provides proper answers. I can get proper JSON from this model, which is good for a production level pipeline.

“The built-in citations feature along with RAG is exceptional. RAG is the top feature this model provides, but along with RAG, the built-in citation feature is very good. RAG responses come with source attribution built in. Users gain more confidence that this model is not hallucinating and is giving proper responses

according to the document. This builds trust with users and the LLM and it is easier to audit which part of the document is being referred to and how accurate the answers are.

“User reactions were very good. Users gained confidence through the built-in citations. Users felt they could improve the system prompt for this model to get better answers because the model was querying the document and after Retrieval-Augmented Generation, this model was able to give answers with citations. This way we gained trust and users were able to gain trust with the model. They were confident that the output is accurate according to the data present in the document. They do not need to assume whether the information is correct. There is no question mark for the trust. I appreciate that part very much about the built-in citations..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions

[Read full review](#) 

# Other Solutions Considered

“Before choosing Cohere Command R, I did not evaluate other options. I cannot share which other options I evaluated before choosing Cohere Command R because there were no other options..”

**Verified user**

Engineer at a construction company with 11-50 employees

[Read full review](#) 

---

“Previously users were facing complaints that the chatbot replied too slowly or the chatbot hallucinated a lot, meaning it gave random answers. Users were complaining, but after using this model, the complaints are very minimal and their support tickets are reduced by 5% to 10%..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions


[Read full review](#) 

“With the legacy system that I used previously, it was a longer, time-consuming process. I had to do all the setup locally and perform RAG tasks to extract data from PDFs. After implementing Cohere Command R, the process became streamlined. The whole process became very fast, and the number of end users increased.

“Many workflows have been removed after using Cohere Command R. For example, when I am using PDFs, I do not have to worry about the extraction part. Even when I am providing some policies, it can extract exact chunks based on user requirements. It also performs well on the citation part..”

**Singh Aman**

Generative AI Engineer at Tata Consultancy

[Read full review](#) 

---

“Other competitors I would have used include OpenAI, which is also very good compared to this model. I have also been using Claude, and Claude models are very good. For RAG, some models include OpenAI 4 models like their GPT-5 pipeline, GPT-5 legacies such as GPT-5.1 and 5.2. The direct match would be the GPT-4o mini model with Cohere Command R because the output and input tokens are the same. Claude's Haiku models are comparable to Cohere Command R, as are Gemini Flash model and Mistral Small model. These are four to five types of models that we have integrated on our platform, so users have the option to use other models if they are not willing to use Cohere Command R.

“The tool call capability of other models was very good compared to Cohere Command R. The output tokens of other models were very large compared to Cohere Command R. With Mistral, web search is very good and fast. There are some benefits of other models too..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions

[Read full review](#) 

# Use Case

“I implemented it myself in my bot, defining what is acceptable based on how people chat and what they might mean when they say certain things. I never actually used that feature for Cohere Command R.

The dataset I am using is just the chat, the user chat, and it is not that big. It is just a few months, and I always clear the chats after a few months. So it is just normal content, nothing extraordinary; I do not think it can be quantified as big data..”

**Collins-Omondi**

Mobile Application Developer at Uamuzi Foundation

[Read full review](#) 

“I have used Cohere Command R mainly for Retrieval-Augmented Generation (RAG) workflows where the model needs to answer questions from enterprise documents rather than relying on its pre-trained knowledge. The primary use cases are document Q&A, chatbot-style assistant, summarizing longer internal content, and providing grounded responses with citations. It is particularly useful when I wanted the model to retrieve the right context from PDFs, similar to what Gemini does in Google Drive. Additionally, it works well with knowledge-based articles, technical documentations, or policy documents to generate concise answers along with source citations, making it clear to the end user where the answers are coming from.

“When I have a document and ask something related to that document, my answer is mostly from the document only. When there is a source citation that indicates it fetched something from that content or paragraphs, it becomes easy to navigate if anything goes wrong. This feature of source citation is among the best. For the context part, when I am asking a longer and repetitive question, it maintains the context of what I was asking, which I find beneficial.

“I would recommend Cohere Command R if your main use case is RAG, document Q&A, and enterprise knowledge assistance, or if you are building any chatbot workflows where citations and grounded answers matter. It strongly fits when you already have a plan to build proper retrieval pipelines..”

**Singh Aman**

Generative AI Engineer at Tata Consultancy

[Read full review](#) 

“On my platform, there were many users who wanted answers from their documents. They had many large-sized documents like PDFs of 25 pages, and some users had PDFs of 150 pages. Using a normal RAG pipeline was very complex for them to get better answers. When we deployed Cohere Command R on our platform, many of our users uploaded their documents and used this model, which gave much better accuracy compared to other models. This was a very good achievement for integrating this model into our platform. The strong tool call use of this model is also very good. With Retrieval-Augmented Generation, this model performs better tool calling as well.

“I can give you some complex examples that users were asking about. Some of my users were NEET students getting their doctor degrees, so their documents had answers related to the human body. There were multiple paragraphs in which the same thing was repeated. They uploaded their book and wanted to learn from it using this model. This model was able to give correct answers from the particular sections. There were mismatches in other models because the same content was referred to somewhere else in another part of the book, but it was not intended to give the answer. This model excelled here. Some of my users had large documents and needed fast answers, with less accuracy being acceptable as long as they got fast responses. They uploaded their documents and received fast answers because they trained their chatbots on the platform. They uploaded FAQ questions and this model was able to give very fast answers according to their document FAQ questions.

“Many of our users came with PDFs and docs needing to build a chatbot for their FAQs on their website. They wanted a chatbot where users could come, create an agent, and load their script on that platform. To train the agent, many of their use cases required providing a chatbot on their website where users could ask FAQ questions. Using Cohere Command R, they were able to upload FAQ questions as a document and this model read the FAQ questions and returned the answer. Many of our chatbots are using Cohere Command R and it returns answers according to user FAQ questions. The accuracy is good. The context window is small, which is good for the FAQ part only because we do not need large context output for FAQ chatbots. The main focus of my company related to Cohere Command R is that we are pitching this model for chatbot selling only.

“After this model release, when we integrated this model on our platform, around 20% of users came to use chatbot. Around 20% of the users were using this FAQ chatbot using Cohere Command R. Previously they were facing complaints that the chatbot replied too slowly or the chatbot hallucinated a lot, meaning it gave random answers. The users were complaining, but after using this model, the complaints are very minimal and their support tickets are reduced by 5% to 10%..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions

[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.

“The documentation of Cohere Command R is very easy. It is a one-page document and everything is written very precisely and crisply. I was able to understand each and everything very easily in a single document. It was very easy to use and integrate. The documents were good..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions

[Read full review](#) 

---

“It takes a few hours, a lot of hours, to deploy Cohere Command R. Not days, but just a lot of hours debugging and dealing with issues when mostly it was on the AWS side, like exposing the API and the static routes. It was just the AWS side of it that took a lot of time, but the model itself was not that complicated..”

**Collins-Omondi**

Mobile Application Developer at Uamuzi Foundation

[Read full review](#) 

# Customer Service and Support

“I have not used customer support for Cohere Command R because I was able to debug everything using their documentation only. I have never reached their customer support. According to the reviews, I would rate them as eight..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions

[Read full review](#) 

---

“I have never interacted with the support team and do not know anyone who has, so I rate it five. I am not saying it is bad; I just have never tried it before, and I do not want to give it a lower score. So I will say five because I hope it is good..”

**Collins-Omondi**

Mobile Application Developer at Umuzi Foundation

[Read full review](#) 

## Other Advice

“There have been significant changes. Developer efforts have been reduced due to streamlined processes. Client and end-user responses have improved with increased trust in responses because of source citations accompanying every paragraph responded to by the model. My overall review rating for Cohere Command R is 9 out of 10..”

**Singh Aman**

Generative AI Engineer at Tata Consultancy

[Read full review](#) 

---

“I do not have anything else to add about how I use Cohere Command R in my projects. I do not know how Cohere Command R has impacted my organization positively.

“My advice to others looking into using Cohere Command R is to try it. I would recommend trying the product. I am giving this review a rating of 9..”

**Verified user**

Engineer at a construction company with 11-50 employees

[Read full review](#) 

---

“I am still using Cohere and maybe [Cloudinary](#).

I work with Cohere Command R occasionally, but not so much.

I am familiar with Cohere Command R, and I just use it as a model. It is pretty

similar to the others that I use, so I cannot really say anything specific about it.

I do not use it every time, just occasionally.

I never tried real-time analysis; I never needed to.

I would rate this review a 7..”

---

**Collins-Omondi**

Mobile Application Developer at Uamuzi Foundation

[Read full review](#) 

“I advise using this model if you are working on RAG-based applications where you have documents and need the model to answer using your document. You should definitely use this model in this case. If you do not want to work on documents, you can choose another model.

“The built-in citations feature along with RAG is exceptional. RAG is the top feature this model provides, but along with RAG, the built-in citation feature is very good. RAG responses come with source attribution built in. Users gain more confidence that this model is not hallucinating and is giving proper responses according to the document. This builds trust with users and the LLM and it is easier to audit which part of the document is being referred to and how accurate the answers are. I would rate this model as seven out of ten..”

**Husain Barwala**

AI Engineer at Walkover Web Solutions

[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](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