




Apache Kafka on Confluent Cloud

Reviews, tips, and advice from real users



Powered by  **PeerSpot**

Contents

Product Recap..... 3 - 5

Valuable Features..... 6 - 12

Other Solutions Considered..... 13 - 14

ROI..... 15 - 16

Use Case..... 17 - 21

Setup..... 22 - 24

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

Other Advice..... 27 - 29

Trends..... 30 - 31

About PeerSpot..... 32 - 33

Product Recap



Apache Kafka on Confluent Cloud

Apache Kafka on Confluent Cloud

Recap

Apache Kafka on Confluent Cloud enables organizations to perform real-time data streaming and processing, integrating user data, and managing large transaction volumes efficiently. Confluent makes it easy to connect and migrate on-premises, multi-cloud, and edge data to AWS. With 120+ pre-built connectors, Confluent allows for real-time streaming pipelines with just a few clicks to Amazon S3, Amazon Redshift, Amazon Kinesis, Amazon Relational Database Service (RDS), Amazon DynamoDB, and more.

Confluent is a cloud-native, data streaming platform that enables you to stream, connect, process, and govern all of your data, allowing you to streamline AI workloads and accelerate data-driven decisions. Confluent's fully managed Apache Flink service available on AWS offers simple, serverless stream processing for building high-quality, reusable data streams. Organizations leverage Apache Kafka on Confluent Cloud for several applications such as capturing change data, throughput or latency sensitive systems, and fostering data migration and microservices. It is instrumental in publishing and managing data across multiple platforms, ensuring both high throughput and seamless log management. Confluent Cloud's features like scalability, efficient transaction management, and cost-effectiveness, combined with seamless integration with multiple data tools, make it a preferred choice for real-time data processing tasks.

What are the key features?

- Scalability: Easily scale up or down based on demand.
- Efficient transaction management: Manage transactions smoothly.
- Cost-effectiveness: Optimize costs without compromising performance.
- Real-time data capture: Capture data in real-time with ease.
- State-saving: Save states efficiently to support various processes.
- Message processing speed: Fast processing of messages ensures quick data handling.
- Topic flexibility: Flexibility in managing topics caters to different data needs.

What benefits and ROI should you look for?

- Seamless integration: Integrate effortlessly with major data tools.

- Streamlined operations: Reduce the number of tools required for data processing and streamline maintenance.
- High throughput: Maintain high throughput across platforms.
- Seamless log management: Effective management of logs ensures operational efficiency.

In specific industries, such as finance, real-time data processing is crucial for handling high volumes of transactions and location tracking. In retail, it integrates user data and manages inventory systems efficiently. IoT applications benefit from Apache Kafka on Confluent Cloud by capturing and processing massive data streams from connected devices in real-time. These are just some of the many applications for Confluent Cloud. The reality is Confluent can be used in any industry across many use cases.

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- ✓ “Confluent helped me to streamline all those logs into one place, and then I was consuming those logs that were produced, which made it very much easier because I know Kafka and using Confluent made it much simpler.”



Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

- ✓ “It's very fast and helps us to create the project, guarantee the message delivery, and the performance.”



FABIO LUIS VELLOSO DA SILVA

Partner at SouJava

- ✓ “Some of the best features with Apache Kafka on Confluent Cloud are streaming and event capabilities, which are important due to scalability and resiliency.”



Verified user

Chief Architect at a financial services firm with 10,001+ employees

- ✓ “Apache Kafka on Confluent Cloud is critical infrastructure for us; without it, our infrastructure costs would increase significantly, potentially amounting to hundreds of thousands of dollars each year, and its real-time capabilities accelerate speed to value and enable new use cases, providing significant business value.”



Anand Venugopal

Global Vice President, Product Strategy & Gtm at NucleusTeg

- ✓ “Apache Kafka on Confluent Cloud is more reliable and frequent to use compared to Apache Kafka.”



Ritik Varshney

Scholar Trainee/Project Engineer at Wipro Limited

- ✓ “The benefits that I have seen from having a real-time architecture include better velocity for developers; instead of developing many of those capabilities in each team, we can rely on Apache Kafka on Confluent Cloud to provide those functionalities we want, and the teams can focus on their own business instead of providing all sorts of APIs and dependencies to other domains, allowing everyone to run faster.”



Verified user

VP Engineering at a tech vendor with 1,001-5,000 employees



“Confluent Cloud handles data volume pretty well.”



Verified user

Integration Solution Architect at a consultancy with 11-50 employees

What users had to say about valuable features:

“Apache Kafka on Confluent Cloud is more reliable and frequent to use compared to Apache Kafka. Its features such as schema registry, cluster linking, error handling, and dead-letter queues provide significant benefits. It also offers enhanced visibility and integration for data streaming, helping clients and customers use it efficiently..”

Ritik Varshney

Scholar Trainee/Project Engineer at Wipro Limited

[Read full review](#)

Apache Kafka on Confluent Cloud is a serverless, multi-cloud SaaS product that eliminates the need for users to manage their own Kafka clusters. It offers numerous connectors to various sources and destinations, facilitating easier integrations. The powerful integration with Flink and Iceberg (Table Flow) enhances functionality. These features are not present in an open-source product. The scalable, reliable service supports multi-cloud data streaming, making it easier for enterprises to connect disparate data sources and destinations.

Anand Venugopal

Global Vice President, Product Strategy & Gtm at NucleusTeg

[Read full review](#) 

“Some of the best features with Apache Kafka on Confluent Cloud are streaming and event capabilities. Streaming and event capabilities are significant. These features are important due to scalability and resiliency.

“In terms of auditing and compliance, Stream Lineage has helped by providing understanding of the entire data domain across the movement and what is happening in the domain and the system..”

Verified user

Chief Architect at a financial services firm with 10,001+ employees

[Read full review](#) 

“The best features Apache Kafka on Confluent Cloud offers would be the connection with various external systems through various languages such as Python and C#. That gave us an edge to use this tool, and we were able to connect to it.

The multi-language support helped my workflow because we did use two kinds of languages, mainly .NET or C# and Python. So we were able to get logs because we had these out-of-the-box connectors available in Confluent. We just had to read the documents, understand how to configure it, how to send the logs, and we were through by just adding a few lines of code in our applications in the respective languages.

Apache Kafka on Confluent Cloud positively impacted my organization because this was just one of our projects we had utilized, while there are many more I was not involved in, so I cannot really speak on them. But this helped a lot with regard to log monitoring and many other streaming use cases where we were able to get all the data in one place and stream it at a quicker rate, and it was much more simpler compared to using plain vanilla Kafka. So Confluent made that difference for us..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“The valuable features with Apache Kafka on Confluent Cloud are the messaging and the asynchronous messages; it's the basic, not advanced usage. It's only to create clusters to receive and send messages.

“The point is the asynchronous messages and the scalability; it is important for us.

“To guarantee the compliance of the architecture and the patterns for the company, to provide scalability, and to guarantee the security to send the messages.

“The Kafka Streams API helps with real-time data transformations and aggregations.

“It's very fast and helps us to create the project, guarantee the message delivery, and the performance. It's a good experience with very impressive processing and a very impressive project and product..”

FABIO LUIS VELLOSO DA SILVA

Partner at SouJava

[Read full review](#) 

“The benefits that I have seen from having a real-time architecture include better velocity for developers. That is the main one. Instead of developing many of those capabilities in each team, we can rely on Apache Kafka on Confluent Cloud to provide those functionalities we want, and the teams can focus on their own business instead of providing all sorts of APIs and dependencies to other domains, allowing everyone to run faster.

“We find that the best features include using the CDC functionality with the connector to take the data from our SQL database and publish it to many consumers. Any changes enable us to easily publish changes about their domain business objects without too much code and work from domain teams. In this way, we can more easily provide a very robust layer of API and events.

“The second use case is easier projection of data. We found that many teams were struggling to create projections and read stores with regular event buses, and Apache Kafka on Confluent Cloud helped us because of all sorts of features, such as the log architecture they have. KSQL also helped us there..”

Verified user

[Read full review](#) 

VP Engineering at a tech vendor with 1,001-5,000 employees

Other Solutions Considered

“We were looking for specific use cases. We compared different Kafka solutions, not necessarily competitors. We have a message bus already. We wanted the log capability, mostly..”

Verified user

VP Engineering at a tech vendor with 1,001-5,000 employees

[Read full review](#) 

“We did not really look at any other options. We were, in fact, given Apache Kafka on Confluent Cloud as the top solution when we were doing our solution designing, so that is the reason we directly started our testing and trials on Confluent Cloud..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“We did initially have Grafana, but the only problem with Grafana was it was limited to certain applications, and ours was not among them. Because of that, we thought of shifting to a unified log monitoring system and started off with having vanilla Kafka installed on our servers. But we found Apache Kafka on Confluent Cloud much more convenient regarding how we would set it up, so to get things done faster, we shifted to Confluent..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“We did not use different streaming analytics before using Apache Kafka on Confluent Cloud.

“We evaluated other options including WebSphere MQ and AMQP, and there was another solution I don't remember..”

FABIO LUIS VELLOSO DA SILVA

Partner at SouJava

[Read full review](#) 

ROI

Real user quotes about their ROI:

Apache Kafka on Confluent Cloud is critical infrastructure for us. Without it, our infrastructure costs would increase significantly, potentially amounting to hundreds of thousands of dollars each year. Its real-time capabilities accelerate speed to value and enable new use cases, providing significant business value.

Anand Venugopal

Global Vice President, Product Strategy & Gtm at NucleusTeg

[Read full review](#) 

“In regards to operational impacts and return on investment, returns depend on the application you deploy and the amount of benefits you are getting, which depends on how many applications you are deploying, what are the sorts of applications, and what are the requirements. If your application has a requirement of eventing, you would have returns to it, but I cannot articulate what sort of dollar value we're talking about..”

Verified user

Chief Architect at a financial services firm with 10,001+ employees

[Read full review](#) 

“Regarding ROI in any capacity, whether it's savings from employees or cloud, the ROI was very significant. Although, specifically with Apache Kafka on Confluent Cloud, it was a bit more challenging to increase adoption because it's very expensive. So we had to pick and choose where we implemented to make sure that ROI is positive.

“I don't remember the exact number because it's been a while since we did a pricing talk, but it was expensive..”

Verified user

[Read full review](#) 

VP Engineering at a tech vendor with 1,001-5,000 employees

“The main return on investment was in the maintenance space because going for Confluent Cloud means you remove all the platform management that you have in terms of these resources that can be allocated to other tasks, where the tool takes basically ownership of all of these. We saw, at the end of the year's end, improvements that were substantial, especially when it comes to the need to resolve issues, as we can deploy the minimum team possible for Confluent because the support model allows for the Confluent team to take ownership of the issue. With AWS, the tool's team supports us, but we have to deploy the right people and take them out of all other initiatives. The most important part is the cost related to platform maintenance and issue resolution..”

Verified user

[Read full review](#) 

Integration Solution Architect at a consultancy with 11-50 employees

Use Case

“In my company, we are not using the tool for analytics and it is more for CDC processes, so we change the capture processes. It is used to extract data from a database and make it available in other parts of our systems or produce events that inform us of data updates..”

Verified user

[Read full review](#) 

Integration Solution Architect at a consultancy with 11-50 employees

I use Apache Kafka on Confluent Cloud as a streaming platform for enterprises to move data in real time from the point of generation to where it needs to be consumed. Use cases for this include point of sale, IoT, financial transactions, and any application that benefits from real-time data processing. My work involves using these solutions for industry verticals and customers in the retail and financial services sectors.

Anand Venugopal

[Read full review](#) 

Global Vice President, Product Strategy & Gtm at NucleusTeq

“We use Apache Kafka on Confluent Cloud for streaming large volumes of data in real-time. It's employed in scenarios such as handling events from various countries and streaming them efficiently for our clients.

We also utilize it for data analytics and in client versions for topic creation, consumer consumption, and ACL provisioning..”

Ritik Varshney

Scholar Trainee/Project Engineer at Wipro Limited

[Read full review](#) 

“We need to send a lot of asynchronous messages in this project, and we use the middleware and Apache Kafka on Confluent Cloud to guarantee asynchronous messaging between the services.

“We use Apache Kafka on Confluent Cloud in the cloud.

“We also use it with AWS..”

FABIO LUIS VELLOSO DA SILVA

Partner at SouJava

[Read full review](#) 

“We find that the best features include using the CDC functionality with the connector to take the data from our SQL database and publish it to many consumers. Any changes enable us to easily publish changes about their domain business objects without too much code and work from domain teams. In this way, we can more easily provide a very robust layer of API and events.

“The second use case is easier projection of data. We found that many teams were struggling to create projections and read stores with regular event buses, and Apache Kafka on Confluent Cloud helped us because of all sorts of features, such as the log architecture they have, and other features. KSQL also helped us there.

“When order is more important, we rely on Apache Kafka on Confluent Cloud..”

Verified user

[Read full review](#) 

VP Engineering at a tech vendor with 1,001-5,000 employees

“I have used Apache Kafka on Confluent Cloud for one of my projects with regard to log monitoring. My main use case for Apache Kafka on Confluent Cloud in that project was mainly streaming of the logs. I wanted to capture logs coming from various interconnected systems into a unified place, so Confluent helped me to streamline all those logs into one place, and then I was consuming those logs that were produced.

Having all my logs unified helped our team a lot because the main challenge we were trying to solve was that in the current scenario we were working on, there was no place where we could view the logs in one place like Grafana or anything similar. It was not available, and for this scenario, we had to use multiple systems in order to check the logs, which could be databases, different applications, and logs for various other APIs. So it was not unified in one place. Now we unified all those logs by producing it to Apache Kafka on Confluent Cloud and then we were consuming it.

It was very much easier because I know Kafka and using Confluent made it much simpler. It was much more easy to understand, grasp, and very well structured with regard to even the JSON response and everything.

I am using Apache Kafka on Confluent Cloud with the Confluent servers itself; we had taken a subscription to Confluent, so it would be the private cloud.

In terms of the development it took us to set this whole thing up using Confluent, we were able to do it at a quicker rate. If we went with the ideal vanilla Kafka, it would require much more manual effort, but here it was easier because of the user interface and the experience, which was mainly very much drag and drop, so we could easily get it done faster.

We did initially have Grafana, but the only problem with Grafana was it was limited to certain applications, and ours was not among them. Because of that, we thought of shifting to a unified log monitoring system and started off with having vanilla Kafka installed on our servers. But we found Apache Kafka on Confluent Cloud much more convenient regarding how we would set it up, so to get things done

faster, we shifted to Confluent..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[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 initial setup is straightforward with provided resources and documentation. I started working with it in the middle stages, not from the initial deployment..”

Ritik Varshney

[Read full review](#) 

Scholar Trainee/Project Engineer at Wipro Limited

“I haven't personally implemented it. The recommendation was made during a consulting program based on a thorough assessment of the tool's capabilities, connectors, and services. However, there hasn't been direct implementation on-site.

.”

DavidBabu

[Read full review](#) 

Senior Architect at a outsourcing company with 501-1,000 employees

“The product's installation phase is pretty straightforward for us since we know how to use it.

Initially, the solution was deployed in a few hours, but after that, we were able to install it quickly..”

Anantharaman Ganesan

Managing Director at Aara Tech Private Limited

[Read full review](#) 

“In terms of the development it took us to set this whole thing up using Confluent, we were able to accomplish it at a quicker rate. If we went with the ideal vanilla Kafka, it would require much more manual effort, but here it was easier because of the user interface and the experience, which was mainly very much drag and drop, so we could easily get it done faster..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“Setting up a new cluster on Apache Kafka on Confluent Cloud is pretty easy. You can click a few options, and there's also integration with other tools to create the necessary resources. However, there are some caveats to be aware of. The type of networking setup you choose can impact your functionality and visibility in the platform. For example, if you have a public cluster, you'll see more metadata information in the console than a more restricted network deployment..”

Verified user

[Read full review](#) 

Tech manager at a tech services company with 1,001-5,000 employees

“I would rate my experience with the initial setup of this product, a seven out of ten where one is difficult and ten is easy.

I have not faced any difficulties or challenges while setting this product up. They have proper documentation, so it's easy to go through it and set things up.

It's the cloud solution, so it's deployed on the cloud in our customers' organizations. And they use Confluent Cloud. .”

Shubham Yadav

[Read full review](#) 

Backend Developer at Mindstix Softwares

Customer Service and Support

The customer support for Apache Kafka on Confluent Cloud is pretty decent. A solid organization supports customers, with many of the original committers and founding team of Apache Kafka involved, reflecting the strong pedigree.

Anand Venugopal

Global Vice President, Product Strategy & Gtm at NucleusTeg

[Read full review](#) 

“I have dealt with the support directly. On a scale of one to 10 with Apache Kafka on Confluent Cloud support, I rate them as good. I would rate them eight if 10 was the best and one was the worst..”

Verified user

Chief Architect at a financial services firm with 10,001+ employees

[Read full review](#) 

“I did interact with the customer support team when I was checking on my credits issued and the cost of my credits utilized. During that time, I had an interaction through email, and it went very well. I was getting prompt responses, and it was nicely handled regarding the support..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“We did not have to ask for much help from the product's support team since we are still exploring the solution. My company mostly looks at forums or blogs, and so we are able to solve any issues associated with the tool. I rate the support an eight out of ten..”

Anantharaman Ganesan

Managing Director at Aara Tech Private Limited

[Read full review](#) 

“We looked at other products, specifically other Kafka providers. We have Apache Kafka and AWS. We looked at self-hosting it, but we wanted Apache Kafka on Confluent Cloud..”

Verified user

VP Engineering at a tech vendor with 1,001-5,000 employees

[Read full review](#) 

Other Advice

“I recommend new users start by going through the [Confluence](#) page and training to learn about Confluent Kafka's features and differences from Apache Kafka..”

Ritik Varshney

Scholar Trainee/Project Engineer at Wipro Limited

[Read full review](#) 

I would recommend Apache Kafka on Confluent Cloud to others, especially those building out a real-time streaming infrastructure or transitioning to real-time business operations from delayed batch processes. They should consider it if they have assets across different clouds. Overall, I rate the solution at eight or eight and a half. It is a stable and steadily growing company with reliable services.

Anand Venugopal

Global Vice President, Product Strategy & Gtm at NucleusTeg

[Read full review](#) 

“My advice to others looking into using Apache Kafka on Confluent Cloud is that it is easier and has a low learning curve. If there is any use case regarding streaming, I would suggest starting off or definitely trying out Apache Kafka on Confluent Cloud and getting the application maybe set up on it, if not evaluating how it returns you the returns. This is a good way to start off because of how it is easier to set up; it saves time, allowing us to utilize the time for other business logic or other things we would need to do. I would rate this solution a 9 out of 10..”

Ashier Fernandes

Lead Software Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“For somebody who's shopping around, looking in this space to decide what to purchase, Apache Kafka on Confluent Cloud is a market leader. It's almost the first choice.

“Going with AWS Apache was also very compelling to us because it's very quick to enable stuff in AWS and try it. I would start with those, but first understand if this is actually what you need. There are other much cheaper solutions that serve other use cases, and sometimes people can mix those and just pick the wrong product.

“Overall, I would rate Apache Kafka on Confluent Cloud a nine out of ten..”

Verified user

VP Engineering at a tech vendor with 1,001-5,000 employees

[Read full review](#) 

“Speaking about data security and privacy requirements, I would say that there are some BAA or legal agreements in the tool. We did not have issues in terms of security or breaches, but before any adoption, with PII or PHI type of data and

before having this data flowing to other clouds or other platforms, the BAA needs to be signed because of IPAC.

Confluent Cloud handles data volume pretty well.

If you are starting to deploy a fully-fledged ETA platform where you do not just have information streaming and go for CDC, and you have some legacy systems that have to communicate on your systems, then I suggest you go for Confluent Cloud.

I rate the tool an eight out of ten..”

Verified user

Integration Solution Architect at a consultancy with 11-50 employees

[Read full review](#) 

“Synchronized messages were very important for us.

“We track performance and throughput as key metrics to evaluate Apache Kafka on Confluent Cloud's handling of event streams.

“The benefits include improving the number of new customers that you can attain, having a more modular system, new functionalities, and high availability.

“The main differences between WebSphere and Apache Kafka on Confluent Cloud are performance, throughput, and the knowledge from the developers in Brazil.

“I have many people working with Apache Kafka on Confluent Cloud in Brazil, and I can send you their contact information if you send me an email.

“I rate Apache Kafka on Confluent Cloud 10 out of 10..”

FABIO LUIS VELLOSO DA SILVA

Partner at SouJava

[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