


aws marketplace

Apache NiFi

Reviews, tips, and advice from real users



Powered by  PeerSpot

Contents

Product Recap..... 3 - 4

Valuable Features..... 5 - 12

Other Solutions Considered..... 13 - 15

ROI..... 16 - 17

Use Case..... 18 - 23

Setup..... 24 - 26

Customer Service and Support..... 27 - 29

Other Advice..... 30 - 34

Trends..... 35 - 36

About PeerSpot..... 37 - 38

Product Recap



Apache NiFi

Apache NiFi Recap

Apache NiFi offers a flexible platform for data orchestration, transformation, and ingestion, catering to both low and high-code customization needs. It streamlines data movement with a powerful visual interface and robust scalability, facilitating seamless integration with diverse data sources.

With Apache NiFi's drag-and-drop capabilities and extensive built-in processors, users can easily simplify complex workflows. Its open-source framework promises cost savings and increased productivity, enabling efficient pipeline development and real-time data handling. While it's valued for data integration and external tool compatibility, there's a need for improvements in logging clarity, local development integration, and cloud-native features.

What are the key features of Apache NiFi?

- **Ease of Use:** A visual interface with drag-and-drop capabilities simplifies operation.
- **Scalability:** Supports both small and large-scale data operations effectively.
- **Integration:** Connects with a wide array of data sources and external tools.
- **Customization:** Offers both low-code and high-code development options.
- **Real-Time Data Handling:** Ensures timely processing and transformation of data.

What benefits can users expect from Apache NiFi?

- **Cost Savings:** Open-source nature helps reduce financial investments.
- **Productivity:** Streamlined processes enhance efficiency and workflow management.
- **Data Handling:** Efficiently manages batch and real-time data for better decision-making.
- **Transformation Options:** Comprehensive tools for flexible data transformation.

In industries like finance, healthcare, and logistics, Apache NiFi is often implemented for data orchestration and transformation tasks, enhancing workflows through integration with tools like Spark and Elasticsearch. It supports data migration and ETL processes, enabling seamless management of large-scale data operations across systems.

Valuable Features

Excerpts from real customer reviews on PeerSpot:



“Apache NiFi is the heart of data ingestion for the entire platform and tremendously makes the platform reliable and useful for customers without crossing any SLAs.”



Yamini Vadlamudi

architect with 51-200 employees



“Apache NiFi has positively impacted my organization as it continually improves functionality and throughput with each iteration over the past three years.”



BJ Tan

Data Engineer at The Kudelski Group



“Apache NiFi has positively impacted the organization by making development really easy, allowing efficient design and development, and enabling code reuse which has reduced the development effort.”



Verified user

Cloud Data Architect at a healthcare company with 10,001+ employees

- ✓ “Apache NiFi provides huge relief for all teams with similar use cases for ETL purposes, and it supports not just ETL but also ELT, allowing us to save significant time.”



Nishant Khandelwal

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

- ✓ “Apache NiFi has positively impacted my organization by definitely bridging the gap between the on-premises and cloud interaction until we find a solution to open the firewall for cloud components to directly interact with on-premises services.”



Debashish Dhar

Manager at a tech company with 10,001+ employees

- ✓ “The ease of use in Apache NiFi has helped my team because anyone can learn how to use it in a short amount of time, so we were able to get a lot of work done.”



Verified user

Data Consultant at a comms service provider with 201-500 employees



“The main benefit of Apache NiFi is that it is an open-source tool, so we don't pay a license for it, and it saves time because we don't have to build connections from the beginning, as it comes with built-in connections.”



AhmedElsaid

Data Management Specialist at a tech vendor with 10,001+ employees

What users had to say about valuable features:

“The best features Apache NiFi offers include the ability to connect to any type of sources, which is a big advantage since most connectors are already available and do not need to be created. In transformation, it has a wide variety of transformations that can be used across big data and any kind of data, including JSON formatted data.

“The connectors used most often include connecting to the Oracle database as the main focus and then connecting to log data, which have greatly benefited the team..”

Verified user

[Read full review](#)

Cloud Data Architect at a healthcare company with 10,001+ employees

“It is easy to integrate Apache NiFi with some data sources, which is a powerful feature in Apache NiFi and it has many connectors to various data sources.

“The main benefit of Apache NiFi is that it is an open-source tool, so we don't pay a license for it. It saves time because we don't have to build connections from the beginning, as it comes with built-in connections.

“In Apache NiFi, the ingestion function is what I really like the most. It is very good at the ingestion part..”

AhmedElsaid

[Read full review](#) 

Data Management Specialist at a tech vendor with 10,001+ employees

“What stood out the most about Apache NiFi was its ease of use when it comes to smaller datasets or when you have multiple streaming data sets that need to be ingested in multiple steps. When you don't have a team with a lot of coding background, Apache NiFi comes in very handy.

The ease of use in Apache NiFi has helped my team because anyone can learn how to use it in a short amount of time, so we were able to get a lot of work done. You don't have to build anything from scratch as you do with Spark or any other tool sets because you already have blocks created within Apache NiFi, so you save time. When we talk about how much time Apache NiFi has saved in development, you save considerable time..”

Verified user

[Read full review](#) 

Data Consultant at a comms service provider with 201-500 employees

“Apache NiFi's best features include its plug-and-play solution, which means you don't need a lot of insight or knowledge to use it. The simplicity and plug-and-play approach of Apache NiFi has helped our team by allowing our customer to scale and build different ingestion flows, which previously needed additional development effort because custom solutions were required; now we just plug and play.

“The features of Apache NiFi, including the integration capabilities to connect to different sources and monitoring with all the observability features, work really well for our team in obtaining the needed information. Apache NiFi has positively impacted our organization by making us a lot more productive; I would say development has significantly improved.

“Development has improved with a reduction in time spent being the main benefit; before we needed a matter of days to create the ingestion flows, but now it only takes a couple of hours to configure..”

Verified user

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

[Read full review](#) 

“The best features Apache NiFi offers include the integration capability because we have use cases wherein we have been using Apache NiFi to integrate through and consume from multiple sources. It can connect to any APIs, NAS drives, and databases, and it can consume streaming data. We can connect to Kafka topics and queues as well through Apache NiFi. First of all, there is flexibility in consuming from multiple sources. We can also easily transfer information from one block to another. We have command and control on the controller services that we can maintain separately so that we do not have to repeat connections and create multiple connections in different processor groups.

“Apache NiFi has positively impacted our organization by significantly helping us streamline the processes. Earlier, each team was creating their own ETL pipelines with no standard being followed. Apache NiFi gave us the opportunity to streamline that process. Each team can request access to Apache NiFi and be onboarded separately based on their needs. We have specifically standardized the pipeline design so that no one can deviate from that. For example, everyone is supposed to use specific variables when enabling the alerting system. We have a different tool called Moogsoft that everyone can onboard to using the InvokeHTTP processor of Apache NiFi to send their alerts to a centralized system which can generate incidents. No one is now working in silos, and there is a specific pattern being followed by everyone. Furthermore, everyone must have Apache NiFi configured so that they can consume from a specific source but eventually push the data to our strategic cloud partner, AWS, first into an AWS bucket which is common for all, and from that bucket, subsequent processing will be done in Snowflake. The standard being followed now once Apache NiFi has been introduced has allowed others to copy-paste successfully created pipelines, saving a lot of time in our overall software development life cycle by reducing redundancy and efforts..”

Nishant Khandelwal

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

[Read full review](#) 

“Apache NiFi offers multiple features. One feature is that it is a no-code platform, so specific code does not need to be written. Direct connectors are available for whatever integrations are needed. If it is Azure, AWS, GCP, or a normal RDBMS, or if connecting to fetch files from a server where files are landing, different types of connectors are available. Integration is simple, flows are created, and the process can run in minutes. Another feature is that if anything needs to be changed or if moving from one flow to a different environment or server is necessary, it is easy to do. Templates can be created, downloaded, and reused as many times as needed. Particular flows can be stored and used multiple times. Email notifications can be configured to trigger if any failures happen, which allows developer teams or the support team to take immediate action. One of the most important features is that if billions and billions of records are coming daily, Apache NiFi ensures that by increasing the configurations and thresholds, that data can be transported, which makes operations easier and allows more things to be accomplished while ensuring SLAs are met.

The most reliable features are the building functionality, the availability of connectors, and the scalability that Apache NiFi offers. These are the features that can be relied on. For example, if billions of records need to be sent from one source to another destination, such as Oracle to Postgres or from a file to an RDBMS, writing a simple Python code requires writing the code and changing the schema. If schema changes occur, it is not easy, as all the code must be reviewed and updated, which takes considerable time for debugging if issues arise. Parallel running of the Python program, multithreading, and all these concepts must be managed, and it is not feasible. The script must be run somewhere and cron must be set up. All these features are bundled and provided in Apache NiFi, which makes operations much easier.

Apache NiFi has tremendously reduced the effort of maintaining multiple different connectors. It makes the transformation and ingestion of billions of records much easier. It is the only platform that is easy for developers to maintain, which makes the entire data platform a very scalable and reliable tool. All data ingestions happen through a single outlet, so if any issues occur, Apache NiFi is the only thing that needs to be addressed, rather than having to go to multiple places and do

multiple things or have debugging become a nightmare if Apache NiFi had not been used. It tremendously makes the platform reliable and useful for customers without crossing any SLAs. Timely ingestions are happening and everything is functioning properly..”

Yamini Vadlamudi

architect with 51-200 employees

[Read full review](#) 

Other Solutions Considered

“We evaluated using Spark for all of our ingestions before choosing Apache NiFi, and at that point, we had to build a framework entirely from scratch. That is why we didn't go there, and we went for Apache NiFi..”

Verified user

Data Consultant at a comms service provider with 201-500 employees

[Read full review](#) 

“We looked at some proprietary solutions and also, evaluated StreamSets. The proprietary solutions were expensive and often didn't suit our use cases. StreamSets didn't have the same level of adoption..”

Verified user

Software Engineer at a consultancy with 1,001-5,000 employees

[Read full review](#) 

“Before Apache NiFi, StreamSets was used, and it was unsatisfactory as it consumed all the server resources and did not release them. After finding Apache NiFi, it became a preferred solution..”

Verified user

Cloud Data Architect at a healthcare company with 10,001+ employees

[Read full review](#) 

“Before choosing Apache NiFi, other options such as StreamSets, Teradata, Informatica Big Data version, and Glue were evaluated, but Glue was not chosen due to its high cost..”

Verified user

[Read full review](#) 

Cloud Data Architect at a healthcare company with 10,001+ employees

I previously used Airflow before switching to Apache NiFi.

I switched from Airflow to Apache NiFi because Airflow is still not widely used within the organization; my experience with Airflow has been with other companies, not with Mastercard. .”

Debashish Dhar

[Read full review](#) 

Manager at a tech company with 10,001+ employees

“We did not use any previous solution before Apache NiFi. Apache NiFi has significantly helped us improve our processes and streamline them, while we are also using SnapLogic for integration purposes in parallel.

“I am not aware of any other solutions that were explored before onboarding Apache NiFi. Being from the analytics team, we were earlier using Tableau Prep for ETL and transformations and Informatica as well, but I am unsure if anything else was explored in the organization prior to using Apache NiFi..”

Nishant Khandelwal

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

[Read full review](#) 

ROI

Real user quotes about their ROI:

“Apache NiFi provides huge relief for all teams with similar use cases for ETL purposes, and it supports not just ETL but also ELT, allowing us to save significant time..”

Nishant Khandelwal

[Read full review](#) 

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

“Apache NiFi provides a good return on investment, although specific cost details are not known since that is not an area of involvement. Feedback has been given to stakeholders that Apache NiFi is beneficial in development and does great work, indicating a positive return on investment..”

Verified user

[Read full review](#) 

Cloud Data Architect at a healthcare company with 10,001+ employees

“I haven't seen a return on investment. There may be return on investment based on the technology and easily moving our workloads onto Apache NiFi from our previous system. However, the tech debt that is created from Apache NiFi is not a good addition. It saves time on development, but then you have to cater to some custom solutions..”

Verified user

[Read full review](#) 

Data Consultant at a comms service provider with 201-500 employees

“I have definitely seen a return on investment through time savings. Working with Apache NiFi allows us to manage it more efficiently, transitioning from spending hours or days resolving issues to requiring much less intervention now. Thanks to improvements on both our side in how we run processes and enhancements to Apache NiFi, we have reduced the time commitment to almost not needing to interact with Apache NiFi except for minor queue-clearance tasks, allowing it to run smoothly. At this point, we have certainly saved hundreds of hours..”

BJ Tan

[Read full review](#) 

Data Engineer at The Kudelski Group

Use Case

“Apache NiFi is used mainly for ETL to get data from multiple sources and then load it into a data lake. For example, data from Ab Initio is extracted and then loaded to an S3 bucket. From the S3 bucket, that data is read again and then loaded into other data layers. Different data layers, such as raw and raw silver, all use Apache NiFi..”

Verified user

[Read full review](#) 

Cloud Data Architect at a healthcare company with 10,001+ employees

“Apache NiFi is mainly used for data ingestion only, functioning as an extraction tool to extract and load data. It is not particularly strong at the transformation part, but it excels at the ingestion part, so we use it as an ingestion tool only.

“For real-time monitoring, we use it for ingestion in real-time. It can be used for either batch ingestion or real-time ingestion, and we use it for both types. It performs well in real-time ingestion scenarios..”

AhmedElsaid

[Read full review](#) 

Data Management Specialist at a tech vendor with 10,001+ employees

“I have been using Apache NiFi for a couple of years now. My experience with Apache NiFi is that I have used it for a customer of ours who used that system to orchestrate ingestion processes from different sources into AWS. My main use case for Apache NiFi is data ingestion, so connecting to sources like APIs or different data lakes to get the data into S3. I manage more than 50 ingestion flows currently; some of them are quite small, dealing with a few megabytes a day, while others are multiple gigabytes a day and those are run daily, so we have a daily refresh of the sources..”

Verified user

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

[Read full review](#) 

“My main use case for Apache NiFi is to use it for basic ETL pipelines before ingesting data into the data lake. For example, if we want to ingest anything into Snowflake, before actually moving it to Snowflake, we do basic data massaging and transformations through Apache NiFi. This may include consuming a file, converting that file from one format to another, breaking it into chunks, and then pushing it to Snowflake.

“A specific example is that we consume all caller data through Apache NiFi by connecting to Kafka topics and consuming those messages. We then convert those messages into JSON format and push those messages to Snowflake by running Snowflake procedures, which eventually ingest the data directly into Snowflake by reading it through an S3 bucket. We also push the actual JSON messages to S3 buckets through Apache NiFi itself..”

Nishant Khandelwal

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

[Read full review](#) 

“I have been using Apache NiFi virtually daily, as it is part of my main responsibility in my current role.

“My main use case for Apache NiFi involves integrating various data sources and performing transformations to load them into mostly our NoSQL database, Elasticsearch, but sometimes into other databases as well.

“For integrating and transforming data, we receive a lot of logs generated with our AWS services that the company wants to collect, particularly for our security team to review those logs and ensure they can conduct their security checks and reviews to confirm there is no abnormal behavior. We use Apache NiFi to capture those logs sent to many S3 buckets, collect those logs, decompress them with Apache NiFi, perform any necessary transformations, and send them to Elasticsearch so that end users, often from the network team or security team, can then use Elasticsearch and Kibana for data analysis.

“My advice for others considering Apache NiFi is that if you are willing to, you can use it on-premises; it offers great customizability. While it is specifically designed for streaming data, it can also accommodate batch data. Moreover, it is useful for various out-of-the-box solutions, including unique uses such as email notifications, showcasing flexibility in data orchestration, ETL, and other applications..”

BJ Tan

Data Engineer at The Kudelski Group

[Read full review](#) 

“Apache NiFi is used to fetch data from different sources and ingest it into different destinations. The entire platform depends on Apache NiFi for data transformation and data movement. Multiple sources such as flat files, Oracle, Postgres are connected, and the main database where data is stored is Postgres and Apache Druid. The UI connects directly to those databases to ensure the data is available at all times. This platform has been built for many enterprise platforms.

Sources include Oracle and files in S3. Apache NiFi creates flows using NiFi processors to fetch data from particular sources such as Oracle. Business data and customer data are available there. Every day there are many incidents, changes, and logging into the system. Data must be fetched daily in two ways: batch streaming and near real-time streaming. Data is fetched twice every day as a batch job. The data is processed after fetching. When connected to the Oracle database, data is received in Avro format. It is then changed into JSON format, and a mapping is used. Whatever destination is required, fields are mapped accordingly. The schema in Postgres is mapped using evaluate JSON path. Whatever JSON field is received is mapped to the respective schema in the database. Records are then ingested into Postgres. Almost daily, 5 million records are received and multiple flows process them. Each flow has a different schedule, which can be every five minutes or twice daily.

Another process retrieves cost data from AWS and Azure. Everyday billing files are received in S3 or Azure blob if Azure is being used. The data is retrieved and transformed with modifications made based on the schema. The data is then sent into Apache Druid using the REST API.

Apache NiFi is the heart of data ingestion for the entire platform. Whatever data movement is happening will happen through Apache NiFi, which is the single data ingestion tool in the platform. Once data lands into the respective destinations using Apache NiFi, multiple transformations, queries, or other operations are performed and projected to the UI. The main gateway for all data from whatever source can be is Apache NiFi. It is a no-code platform where multiple flows can be created, which reduces the time spent writing code and integrating into different connectors. Connectors are already available in Apache NiFi, so flows just need to

be created, which makes the process easy and stable..”

Yamini Vadlamudi

architect with 51-200 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 was straightforward. We can deploy the tool easily on a single node. It won’t take much time. If it is a multi-node cluster, it will take two to three hours..”

MoulaliNaguri

Project Engineer at Wipro Limited

[Read full review](#) 

“I tried to install the tool on my work laptop, and while it worked initially, it started to run slowly after some time. The department that handles the company's databases uses Apache NiFi on proper servers. I tried using it on my laptop to see if it worked, but it ran very slowly and consumed many resources from my machine..”

Teodor Muraru

Developer at Emag

[Read full review](#) 

“Apache NiFi is deployed in our organization using public cloud infrastructure. I purchased Apache NiFi through the AWS Marketplace. There are no improvements needed for Apache NiFi deployment on AWS, as the experience with that was pretty smooth..”

Verified user

[Read full review](#) 

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

“The initial setup is fast, especially for communication stabilization. Although the product is open source, it functions as a cluster. For single-node environments, installation is simple. For company-wide or enterprise-level clusters, the initial stages may present issues with authentication and access. Stabilization, such as port communication, may not be immediately effective..”

Bharghava Raghavendra Beesa

[Read full review](#) 

Senior Developer at Infosys

“The initial setup is very easy, especially for users who are familiar with EDL or EMT.

NiFi is one of the easiest tools on the market to learn and use. It is also a quick-win solution, which is good for first-time users who are developing data pipelines for EMT. NiFi makes it easy to track and trace the status of your pipelines, so you can be sure that they are working properly..”

Arjun Pandey

[Read full review](#) 

Engineering Lead- Cloud and Platform Architecture at a financial services firm with 1,001-5,000 employees

“The initial setup is very easy. I would rate my experience with the initial setup a ten out of ten, where one point is difficult, and ten points are easy.

But if you want its custom mode and control, it's five out of ten.

For the initial setup, if you configure to custom mode, it's five points. But if you use its single-mode configuration and installation, it's ten..”

SabinaZeynalova

[Read full review](#) 

Data Engineer Team Lead at Unibank

Customer Service and Support

The technical support from the official Apache team is rated a three out of ten. Issues often require self-resolution or community help, as the support isn't effectively managed.

Danuphan Suwanwong

Head of Data Engineering and AI Engineering at Coraline

[Read full review](#) 

“I have had to reach out for help with Apache NiFi, but since I get Apache NiFi from Cloudera, they have their premium support, so it is not that big an issue..”

Verified user

Data Consultant at a comms service provider with 201-500 employees

[Read full review](#) 

“Initially, customer support was contacted more often, but after understanding Apache NiFi better, not many issues have been faced. The customer support is really good, and they are helpful whenever concerns are posted, responding immediately..”

Verified user

Cloud Data Architect at a healthcare company with 10,001+ employees

[Read full review](#) 

“Customer support for Apache NiFi has been excellent, with minimal response times whenever we raise cases that cannot be directly addressed by logs. The support team has consistently provided great assistance with processor failures and helped us create ad-hoc processors as needed..”

Nishant Khandelwal

[Read full review](#) 

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

“We contacted technical support by phone once and they were helpful, but there was some delay with their response. We would have preferred a quicker response time.

After a bit of study, we can support most of our needs so don't really need technical support. .”

Bruno da Silva

[Read full review](#) 

Senior Manager at Timestamp, SA

“I have not needed Apache support. With my seniority, I have not used anything that requires their support.

It's a good thing to be able to go back in the documentation and the information is enough to get you going.

Juniors or people who have less experience may have a lot of different questions, but I have not tested their support..”

Reza Gholami

Senior Technology Architect at a tech services company with 10,001+ employees

[Read full review](#) 

Other Advice

“The process and cost consumption, I believe, is actually not that good when working with Apache NiFi. Apache NiFi is one of the most expensive tools we use. I do believe it is an obsolete technology and the world is moving away from it. I have given this review a rating of five out of ten..”

Verified user

Data Consultant at a comms service provider with 201-500 employees

[Read full review](#) 

“We currently do not work with [TIBCO Jaspersoft](#). I can talk about Apache NiFi and I have experience with the product. We use a reporting solution called [SAP BusinessObjects](#). We use solutions like IBM DataStage, Apache NiFi, [Apache Kafka](#), and [Apache Spark](#). I am not working with SkyWalking from Apache. From other vendors, I don't know if there is a competitor of Apache in this specific role. I would rate this review a seven overall..”

AhmedElsaid

Data Management Specialist at a tech vendor with 10,001+ employees

[Read full review](#) 

“The reason I rate Apache NiFi a seven is that most development folks are afraid to start using Apache NiFi because, to begin with, it does not always directly make sense. For example, there are other integration tools such as [SnapLogic](#) where you can simply search for a specific processor, but that is not the case with Apache NiFi. You need some basic understanding of which processors are there before you can fetch what you need. Better UI design should allow newcomers to search using

relevant keywords, such as API, to retrieve appropriate processors, which currently is not happening, requiring some understanding of Apache NiFi. The other challenge I mentioned is having better logging, especially for processor-related logs to help newcomers navigate effectively.

“I think Apache NiFi is a great tool that one should definitely explore. It is essential to perform basic checks regarding requirements, but if someone is looking for ETL and ELT functionalities needing to connect to CSVs, JSONs, Excels, and databases, they can onboard to Apache NiFi. It offers great connectivity for consuming or pushing data through queues and cloud workloads. Overall, it is an excellent product, especially for basic data massaging and processing before pushing data to Snowflake or creating reports. I rate Apache NiFi a seven out of ten..”

Nishant Khandelwal

Team Lead Technical Specialist (Production Support) at a recreational facilities/services company with 11-50 employees

[Read full review](#) 

“The customizability of Apache NiFi helps even with unique use cases, as I mentioned before, given that Apache NiFi can be used in this capacity. While there are better applications or software options available, when you are trying to keep it simple and finding ways to utilize a couple of processors for a unique solution, you can do that in Apache NiFi. For example, we have several notification-type pipelines we have built in Apache NiFi, such as reading from a SQL database to identify users who have not completed training and then sending them an email reminder to complete that training. We have that running regularly, week by week. Another instance involves a processing data flow that scans for specific data found in logs, which triggers an email notification to the relevant team letting them know that a unique identifier has appeared, allowing them to handle the situation.

“I encountered some odd cases such as increasing concurrent threads on a processor, which should work similarly to copying several processors, yet functional throughput varies. It seems that using a distributed processor yields

better throughput than just increasing the concurrent threads on one processor, which has been odd but is a workaround we had to adopt to boost throughput. Resolving such quirks could elevate the rating further.

“I rate Apache NiFi an eight out of ten. I choose eight because, as open-source software, there is always room for improvement, but the tradeoff between learning how to use the software and the savings it provides, along with its customizability, ranks it pretty high. It is effective for what it does and continues to improve, so it could score higher if there are significant enhancements in custom-built processors and ongoing improvements in functionality..”

BJ Tan

Data Engineer at The Kudelski Group

[Read full review](#) 

“The first benefit is the time saved by moving to Apache NiFi. Generally, writing a Python program and creating different connectors using a Python program may take three weeks to one week for a developer to create, debug, and deploy. When using Apache NiFi, the time was reduced by almost half. Flows can be created in two days a week, whereas with Python it took almost one to one and a half weeks to complete a single flow. In terms of debugging, Apache NiFi is the only platform where errors can be easily debugged. A bulletin board is available where the exact issue can be viewed and traced back to see where it went wrong. These tasks can be completed quickly, possibly in 30 to 50 minutes to debug an issue. Whereas with Python, if working with a single program level and creating multiple files and multiple jobs, it is very difficult to go back and see exactly where the issue is because tracing from the start to find the issue is necessary, and then every line of code must be debugged. With Apache NiFi, the issue location is immediately identified, and focus can be placed on that particular processor. SLAs have been improved because if an SLA requires completing a job within 30 minutes, using a native programming language or normal jobs requires refinement and optimization for very large transformations, such as ingesting 5 billion records

every day, which is not feasible with a normal Python program. Apache NiFi provides the infrastructure where an unlimited number of records, such as billions of records, can be processed in minutes, which ensures that SLAs have never been missed or crossed. SLAs are continuously met, and if additional processing is needed, threads can be increased and little configurations in Apache NiFi can be adjusted, which ensures SLA compliance. These are the metrics that have been seen and benefited from.

If an organization is looking for a unified tool to ingest data at very large scale and wants both batch and real-time streaming with one unified data tool, Apache NiFi is recommended. If a code-free approach is preferred and spending much time writing code or debugging code should be avoided, Apache NiFi is an excellent choice where data flows can be created simply through clicks instead of using other particular tools. This review has been given a rating of 8 out of 10..”

Yamini Vadlamudi

architect with 51-200 employees

[Read full review](#) 

“To train or onboard new team members to use Apache NiFi, the code has been modularized and processes have been documented really well, so when new team members are onboarded, they are asked to review that documentation to understand the processes and the modules that have been created. In a couple of days, once they go through all that material, they are up to speed.

“In terms of flexibility and ease of use, Apache NiFi is more open compared to other ETL tools that have been used, such as Informatica and [Teradata](#). It is open source with many contributors and can handle various data sources, including log data, structured, unstructured, and semi-structured data, unlike traditional ETL tools.

“Tools such as Prometheus and [Grafana](#) are sometimes used to keep an eye on Apache NiFi server, and DataDog is also used along with it. Scaling Apache NiFi

workloads is managed through auto-scaling.

“To handle data security and compliance when using Apache NiFi, LDAP authentication is utilized, all clusters and nodes are kerberized, and single sign-on is used to authenticate. In transit, SSL encryption is used, and at rest, AES encryption is used, which is more than enough for the needs.

“Apache NiFi is kept up to date by keeping an eye on new features that have been released, discussing them internally to assess if they need to be incorporated into development. If there are any gaps in the current version, an upgrade to the new version will be attempted.

“Apache NiFi is a pretty good tool that meets most ETL needs, and in terms of performance and security, it is really good. After using it for quite some time without any issues, it is recommended as the number one tool for ETL. The overall review rating for Apache NiFi is 8 out of 10..”

Verified user

Cloud Data Architect at a healthcare company with 10,001+ employees

[Read full review](#) 


Top Industries

by visitors reading reviews




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 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