

aws marketplace

Dremio

Reviews, tips, and advice from real users



Powered by  PeerSpot



Contents

- Product Recap..... 3 - 4
- Valuable Features..... 5 - 10
- Other Solutions Considered..... 11
- ROI..... 12
- Use Case..... 13 - 16
- Setup..... 17 - 18
- Customer Service and Support..... 19 - 20
- Other Advice..... 21 - 23
- Trends..... 24 - 25
- About PeerSpot..... 26 - 27

Product Recap



Dremio

Dremio Recap

Dremio offers a comprehensive platform for data warehousing and data engineering, integrating seamlessly with data storage systems like Amazon S3 and Azure. Its main features include scalability, query federation, and data reflection.

Dremio's core strength lies in its ability to function as a robust data lake query engine and data warehousing solution. It facilitates the creation of complex queries with ease, thanks to its support for Apache Airflow and query federation across endpoints. Despite challenges with Delta connector support, complex query execution, and expensive licensing, users find it valuable for managing ad-hoc queries and financial data analytics. The platform aids in SQL table management and BI traffic visualization while reducing storage costs and resolving storage conflicts typical in traditional data warehouses.

What are Dremio's most valuable features?

- **Native Error Interfaces:** Simplifies error detection and troubleshooting
- **Integration with Apache Airflow:** Enhances workflow automation capabilities
- **Scalability:** Effortlessly handles growing data demands
- **Role-Based Access Management:** Provides secure data access control
- **Data Reflection:** Optimizes query performance automatically
- **Query Federation:** Supports queries across multiple endpoints seamlessly
- **Data Cataloging and Virtualization:** Offers streamlined data organization

What benefits and ROI should users look for?

- **Reduced Data Storage Costs:** Efficiently manages data to minimize storage expenses
- **Improved Query Performance:** Enhances processing speed for large datasets
- **Flexibility in Integration:** Easily connects with storage solutions like Oracle and MySQL
- **Enhanced Data Management:** Supports comprehensive data analysis workflows

Dremio is primarily implemented in industries requiring extensive data engineering and analytics, including finance and technology. Companies use it for constructing data frameworks, efficiently processing financial analytics, and visualizing BI traffic. It acts as a viable alternative to AWS Glue and Apache Hive, integrating seamlessly with multiple databases, including Oracle and MySQL, offering robust solutions for data-driven strategies. Despite some challenges, its ability to reduce data storage costs and manage complex queries makes it a favorable choice among enterprise users.

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- ✓ “The first feature that stands out for me in Dremio is the federated type of query, which allows the possibility to use multiple endpoints without worrying about writing custom SQL that runs only for SQL Server or for Postgres and Redshift.”



Corrr Moray

SR BI developer at BRQ Digital Solutions

- ✓ “Dremio has positively impacted my organization by helping us create a single source of truth, a singular data warehouse where we can have access to all of the data sets.”



Syedlsmail

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

- ✓ “Dremio has positively impacted my organization as nowadays we are connected to multiple databases from multiple environments, multiple APIs, and applications, and Dremio organizes everything in an amazing way for me.”



Joao Silveira

Data Analyst at a insurance company with 501-1,000 employees

- ✔ “It's almost similar, yet it's better than Starburst in spinning up or connecting to the new source since it's on SaaS.”



KamleshPant

Senior Software Architect at USEReady

- ✔ “Dremio is very easy to use for building queries.”



Vu Thang

Investment Banking Associate at Vietcap

- ✔ “We primarily use Dremio to create a data framework and a data queue.”



Sachin Shukre

Sr Manager at a transportation company with 10,001+ employees

- ✔ “Everyone uses Dremio in my company; some use it only for the analytics function.”



Victor A.

Senior Software Engineer Technical Lead at Apple

What users had to say about valuable features:

“Dremio is very easy to use for building queries. It's easy to connect Dremio to various databases and data sources like Oracle and MySQL. It is also very flexible, providing us with scalability and integration capabilities effortlessly..”

Vu Thang

Investment Banking Associate at Vietcap

[Read full review](#) 

“The most valuable feature is that you can generate refresh reflections and create your visuals (VDS) because it makes it easier to monitor day-to-day data structure. I can use Dremio to create a visual table without impacting the original by creating an opportunity on my own. I can work on my videos and create reflections, and Dremio allows me to use those reflections. I can know the health of my tables, whether healthy or unhealthy, daily. It is one of the most valuable features of Dremio..”

Victor A.

Senior Software Engineer Technical Lead at Apple

[Read full review](#) 

“It's almost similar, yet it's better than Starburst in spinning up or connecting to the new source since it's on SaaS. It is a similar experience between the based application and cloud-based application. You just get the source, connect the data, get visualization, get connected, and do whatever you want.

They say data reflection is one way where they do the caching and all that. Starburst also does the caching. In Starburst, you have a data product. Here, the data product comes from a reflection perspective. They are working on a columnar memory map, columnar computation. That will have some edge in computation..”

KamleshPant

Senior Software Architect at USEReady

[Read full review](#) 

“The best features Dremio offers, in my opinion, probably include the web environment.

What stands out to me about the web environment is both the interface and the speed; it helps a lot.

In terms of features, querying is very easy and the integration with multiple dashboards is amazing.

Dremio has positively impacted my organization as nowadays we are connected to multiple databases from multiple environments, multiple APIs, and applications, and Dremio organizes everything in an amazing way for me.

A specific outcome that shows how Dremio has helped my organization is definitely better data access..”

Joao Silveira

Data Analyst at a insurance company with 501-1,000 employees

[Read full review](#) 

“The best features Dremio offers include having a single system where we can manage all of our data cataloging and visualization or virtualization.

“The interface is a plus over the traditional warehousing solutions, which makes it easier to work with Dremio compared to other solutions I've used.

“Having everything under one system and an easier-to-work-with interface, along with having API integrations, adds significant value to working with Dremio.

“Dremio has positively impacted my organization by helping us create a single source of truth, a singular data warehouse where we can have access to all of the data sets. The fact that Dremio has a clear role-based access management system helps us significantly, as we can have roles segregating all of the data, while users with the appropriate roles can access everything..”

SyedIsmail

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

[Read full review](#) 

“The first feature that stands out for me in Dremio is the federated type of query, which allows the possibility to use multiple endpoints without worrying about writing custom SQL that runs only for SQL Server or for Postgres and Redshift. We build a SQL on Dremio standard, and then Dremio goes over and rewrites the whole lineage to make it possible to use multiple sources of data.

The feature impacts our daily work by reducing complexity because we don't need to worry about where the data comes from. We are slightly migrating to Snowflake, so Dremio is currently being replaced because Snowflake is a more robust platform, but we are kind of happy with the work Dremio does.

I think the reduction of complexity is a positive impact that Dremio has had on my organization. The main thing is that it is a place where you can write down simple SQLs and see the lineage, the way it integrates together; this is the best part of it..”

Corrr Moray

SR BI developer at BRQ Digital Solutions

[Read full review](#) 

Other Solutions Considered

“I didn't get to choose the solution I used before Dremio, which was SQL Server from Microsoft. Nowadays I see that SQL Server solution is much more simple, and Dremio helps me a lot to move beyond that simplicity..”

Joao Silveira

Data Analyst at a insurance company with 501-1,000 employees

[Read full review](#) 

“Before Dremio, we used Hive and AWS Glue, and we wanted to shift to an open-source version that is newer than Hive and provides flexibility in moving between cloud and on-premises..”

Syedlsmail

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

[Read full review](#) 

“I probably don't remember all the options we evaluated, but we talked a lot about Databricks and Redshift, and Dremio was the best choice. Also, we don't use Tableau Prep because we use Dremio, which is an important thing..”

Corrr Moray

SR BI developer at BRQ Digital Solutions

[Read full review](#) 

ROI

Real user quotes about their ROI:

“I cannot share the exact metric itself, but Dremio surely saves time, reduces costs, and all those things because we don't have to worry so much about the infrastructure to make the different tools communicate. Dremio was a solution that made it possible for us to have different data sources with only one registration on our data lake..”

Corrr Moray

SR BI developer at BRQ Digital Solutions

[Read full review](#) 

Use Case

“I use Dremio for proof of concept purposes. I haven't used it in a real-time project, however, I explore Dremio as a data virtualization application in the ecosystem. It is relatively new, possibly a one-year or two-year-old system..”

KamleshPant

Senior Software Architect at USEReady

[Read full review](#) 

“I can visualize traffic from BI and Tableau on the same page and have my tables and schema on the same page. The data link comprises everything. If I want one structure, I connect it to a big table in the hive and the data team that could read my SQL work on my tables, schemas, table structures and everything all in one place. Dremio is as good as any other Presto engine..”

Victor A.

Senior Software Engineer Technical Lead at Apple

[Read full review](#) 

“As a BI developer, my main use case for Dremio involves registering data zones and also doing inquiries, views, and making it persistent to use on BI tools, like Tableau reports. In a project, we create data zones, typically one that we have in our company to receive files. After doing the ingestion inside Dremio, what we do is set up permissions, and after setting up the permissions, the refresh session allows us to go over and start building the views and all those things to be possible to connect via JDBC driver and consume on Dremio dashboard..”

Corrr Moray

SR BI developer at BRQ Digital Solutions

[Read full review](#) 

“We use Dremio for financial data analytics and as a data lake. We connect Dremio with Oracle, Docker, MySQL, and utilize it for Power BI.

Additionally, we use it to process data from MongoDB, although we face occasional challenges with NoSQL integration..”

Vu Thang

Investment Banking Associate at Vietcap

[Read full review](#) 

“I have been using Dremio for a year and a half. My main use case for Dremio is that I am able to access multiple databases and I can easily and quickly connect Dremio with my dashboards.

In my recent project, my databases are full of information about our insurance operation, and I monitor my KPIs by connecting Dremio to my dashboards. I have multiple KPIs that I have to stay tuned to, and I need to access that very quickly and in an easy way, which is how Dremio helps me.

About my use case with Dremio, I find databases that I don't know very much about, and in that web environment of Dremio, I am able to scroll and navigate through the databases by searching for an example or a column name, so it is much easier to navigate through with that kind of help..”

Joao Silveira

Data Analyst at a insurance company with 501-1,000 employees

[Read full review](#) 

“I have been using Dremio on and off as a data warehouse for the past three years.

“My main use case for Dremio is that we use it as a logical data warehouse where we use Dremio with VDSs as an alternative to AWS Glue or Apache Hive. As we are working with our ETL at the end of all of it, after the data types and everything have been cast, we make that available on Dremio as VDS and then we move on to our further data warehousing schemes within Dremio.

“We use Dremio enterprise-wide now, and the key use case has been reducing our costs when it comes to data storage.

“Our main use case for Dremio is as a data warehouse, and the challenge that it helped us solve is that physical data warehouses such as Redshift have storage and hardware upscaling conflicts. Dremio helps us decouple those and lets us catalog more. We can manage everything under one system..”

Syedlsmail

Senior Consultant - Data Analytics at a comms service provider with 201-500 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 installation process was quite smooth and didn't present any issues.

We currently have Dremio on the cloud. For proof of concept (POC) purposes, we are using it on-premises..”

Sachin Shukre

[Read full review](#) 

Sr Manager at a transportation company with 10,001+ employees

“The initial setup is pretty straightforward. I've set it up in two different ways. The first was a three-node configuration. I used virtual machines to set up three different nodes and cluster data. That worked fairly well. I also created containerized versions using Kubernetes. It was much easier to do that. Dremio also has cloud versions where you don't need to do anything. You can go to the cloud and consume it..”

MikeWalker

[Read full review](#) 

Founder/ CEO at Morphogen Data Inc.


“The complexity of the setup depends on the use case. If you are working with a cluster that you plan to shut down after a week, the setup is simple. For more complicated scenarios where you need to connect parquet table, it is more complicated and the documentation to support this setup is not comprehensive.

Experienced DevOps teams need to spend at least three to five days to understand the level of ability between multiple masters and how the scheduling works.

I tried to integrate my Dremio metrics and logs to DataDog and it wasn't simple..”

Nick Grigoriev

Senior Data Engineer at AppsFlyer


[Read full review](#) 

Customer Service and Support

“I have used the community forum to get most of my answers for Dremio. This is best suited for simple queries but does not work as well for advanced questions that need input from a developer. .”

Nick Grigoriev

Senior Data Engineer at AppsFlyer

[Read full review](#) 

“We have had to reach out for customer support many times, and they respond, so they are pretty supportive about some long-term issues. Many times it's with vendors where there is a missing feature or ongoing problem that they say will be fixed in the next release, but this happens not only with Dremio but with almost every vendor we have..”

Corrr Moray

SR BI developer at BRQ Digital Solutions

[Read full review](#) 

“I rate Dremio eight out of 10. Dremio focuses on enterprise customers, so they have a solid support organization. They have professional services and full support. They have a great customer-centered team, but there are some challenges. For example, they don't have enough staff. As more customers buy Dremio, they need more integrators to help them get started. They can't do it all themselves..”

MikeWalker

Founder/ CEO at Morphogen Data Inc.

[Read full review](#) 

Other Advice

“My advice for others looking into using Dremio is to explore the web environment and stay sharp with the ODBC connectors and drivers, as that is very important. I would rate this product a 9..”

Joao Silveira

Data Analyst at a insurance company with 501-1,000 employees

[Read full review](#) 

“Dremio is very flexible and easy to use, making it very suitable for our team. I would recommend Dremio for similar use cases due to its flexibility.

On a scale of one to ten, I would rate Dremio at eight..”

Vu Thang

Investment Banking Associate at Vietcap

[Read full review](#) 

“I suggest that you give it a try and see for yourself. Dremio is as fast as Presto. For example, ten billion rows can return in less than five seconds.

I would rate it a 10 out of 10. .”

Victor A.

Senior Software Engineer Technical Lead at Apple

[Read full review](#) 

“My advice to others looking into using Dremio is that it is a great tool because it keeps all your efforts together. It's a good thing to have if you want to have a unified catalog or metadata or something like this, so this is certainly a good tool to consider. I would rate this review an eight out of ten overall..”

Corrr Moray

SR BI developer at BRQ Digital Solutions

[Read full review](#) 

“Regarding features, I'm not sure if they have all the tools like data governance, data quality, and data lineage integrated. If not, they need to build those tools as well to check the data quality and lineage. Data discovery is there. Connectivity-wise, Starburst is way better, however, Dremio might have a better computing path, possibly delivering data faster than Starburst. No direct comparison can be made, so I cannot comment further.

Overall, you can rate it as eight out of ten..”

KamleshPant

Senior Software Architect at USEReady

[Read full review](#) 

“I would advise others looking into using Dremio to study the tool beforehand. Dremio has several different offerings, and the best way to get into it is to use the open-source on-premises version to experiment with it. To extract the complete power of the platform, organizations should educate themselves with the complete information and compare it with other solutions since a single solution cannot fit everywhere. Educating the team before adopting a technology is better than just adopting a suggested package. I rate Dremio 8 out of 10..”

SyedIsmail

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

[Read full review](#) 

Top Industries

by visitors reading reviews

Financial Services Firm

26%

Computer Software Company

8%

Manufacturing Company


6%

Healthcare Company


6%

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