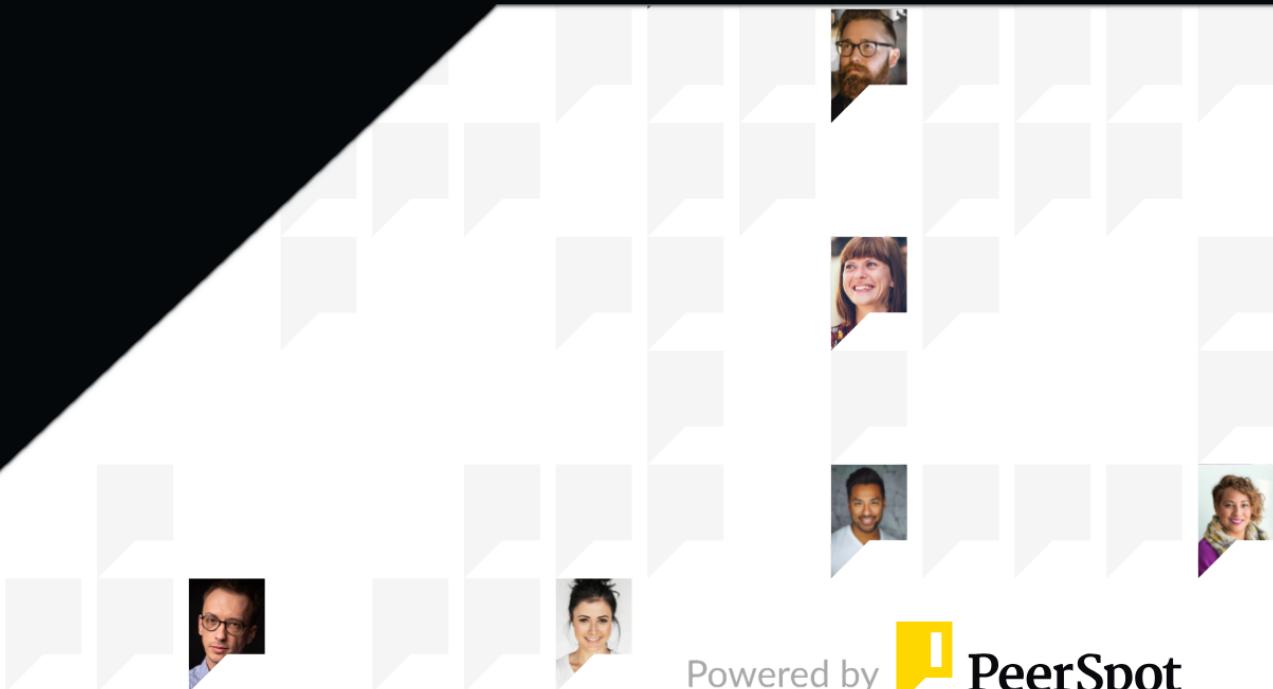




Dremio

Reviews, tips, and advice from real users



Powered by  PeerSpot

Contents

Product Recap	3 - 5
Valuable Features	6 - 11
Use Case	12 - 14
Setup	15 - 16
Customer Service and Support	17
Other Advice	18 - 20
Trends	21 - 22
About PeerSpot	23 - 24

Product Recap



Dremio

Dremio Recap

Dremio is a data analytics platform designed to simplify and expedite the data analysis process by enabling direct querying across multiple data sources without the need for data replication. This solution stands out due to its approach to data lake transformation, offering tools that allow users to access and query data stored in various formats and locations as if it were all in a single relational database.

At its core, Dremio facilitates a more streamlined data management experience. It integrates easily with existing data lakes, allowing organizations to continue using their storage of choice, such as AWS S3, Microsoft ADLS, or Hadoop, without data migration. Dremio supports SQL queries, which means it seamlessly integrates with familiar BI tools and data science frameworks, enhancing user accessibility and reducing the learning curve typically associated with adopting new data technologies.

What Are Dremio's Key Features?

- **Data Reflections:** Reduces query times by creating optimized representations of source data, which can accelerate performance without the complexity of traditional data warehousing solutions.
- **Semantic Layer:** Allows users to define business metrics and dimensions centrally, ensuring consistency and governance across all analytics tools.
- **Built-in Security Features:** Provides robust security measures, including column- and row-level security, ensuring compliance with data governance and privacy standards.
- **Support for Multiple Data Formats and Sources:** Enables querying directly against a variety of data formats (Parquet, JSON, etc.) and sources without the need for conversion or replication.

What Benefits Should Users Expect?

When evaluating Dremio, potential users should look for feedback on its query performance, especially in environments with large and complex data sets. Reviews might highlight the efficiency gains from using Dremio's data reflections and its ability to integrate with existing BI tools without significant changes to underlying data structures. Also, check how other users evaluate its ease of deployment and scalability, particularly in hybrid and cloud environments.

How is Dremio Implemented Across Different Industries?

Dremio is widely applicable across various industries, including finance, healthcare, and retail, where organizations benefit from rapid, on-demand access to large volumes of data spread across disparate systems. For instance, in healthcare, Dremio can be used to analyze patient outcomes across different data repositories, improving treatment strategies and operational

efficiencies.

What About Dremio's Pricing, Licensing, and Support?

Dremio offers a flexible pricing model that caters to different sizes and types of businesses, including a free community version for smaller teams and proof-of-concept projects. Their enterprise version is subscription-based, with pricing varying based on the deployment scale and support needs. Customer support is comprehensive, featuring dedicated assistance, online resources, and community support.

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- ✓ “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

- ✓ “Dremio enables you to manage changes more effectively than any other data warehouse platform. There are two things that come into play. One is data lineage. If you are looking at data in Dremio, you may want to know the source and what happened to it along the way or how it may have been transformed in the data pipeline to get to the point where you're consuming it.”



Mike Walker

Founder/ CEO at Morphogen Data Inc.

- ✓ “Dremio gives you the ability to create services which do not require additional resources and sterilization.”



Nick Grigoriev

Senior Data Engineer at AppsFlyer

- ✓ “Dremio allows querying the files I have on my block storage or object storage.”



Verified user

Security Data Engineer at a pharma/biotech company with 5,001-10,000 employees

What users had to say about valuable features:

“We're still in the exploration phase with Dremio, so it's a bit early to determine its most valuable feature. We're currently deploying it across different departments for various use cases and learning from these internal applications..”

Sachin Shukre

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

[Read full review](#) 

“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](#) 

“Dremio allows querying the files I have on my block storage or object storage. The solution gives me a place where I can play around with the data virtually by creating VDSs or PDSs. Dremio works just like a data warehouse on top of my data lake, which is interesting..”

Verified user

Security Data Engineer at a pharma/biotech company with 5,001-10,000 employees

[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](#) 

“Dremio enables you to manage changes more effectively than any other data warehouse platform. There are two things that come into play. One is data lineage. If you are looking at data in Dremio, you may want to know the source and what happened to it along the way or how it may have been transformed in the data pipeline to get to the point where you're consuming it.

There's another thing called data providence. They're tied together. Data providence allows you to go back and recreate the data at any particular point in time. It's extremely important for compliance and governance issues because data changes all the time. How did it change? What was it three days or months ago? You may have made some decisions based on data that was three months old, so you might need to revisit those.

It's essential for things like machine learning and deep learning, where you are generating AI models off data. When the model stops working or doesn't work as expected, you need to figure out why. You have to go back and adjust the datasets used to train the model. We do that through an open-source project called Nessie, which is their basis for providing data lineage and data provenance capabilities. It's super powerful.

Arrow is another open-source project for storing data in memory and performing data query operations. Data sits on a disk in one format. If you want to do anything with data, you have to load it into your computer and put it into memory so you can work with it. Arrow provides a format in memory that enables the whole library to perform various operations on that data. Every vendor has its own way of representing data in memory.

They've latched onto an industry standard and developed it so it's open. Now people can use the exact same format in memory to do operations and use the library set to perform functions on data. New developers can decide if they want to develop their own memory format or use one that's already there. Data transfer is a massive problem when you're working with large datasets, doing advanced analytics, and trying to train machine learning or deep learning models. What happens often is companies downsample their data sets to do

training on models because transferring and managing data on a deep learning or machine learning platform is too much..”

MikeWalker

Founder/ CEO at Morphogen Data Inc.

[Read full review](#) 

Use Case

“We have been using it to build one of our frameworks. We primarily use Dremio to create a data framework and a data queue. It's being used in combination with DBT and Databricks..”

Sachin Shukre

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

[Read full review](#) ↗

“Dremio is a platform that enables you to perform high-performance queries from a data lake. It helps you manage that data in a sophisticated way. The use cases are broad, but it allows you to make extremely good use of data in a data lake. It gives you data warehouse capabilities with data lake data..”

MikeWalker

Founder/ CEO at Morphogen Data Inc.

[Read full review](#) ↗

“We were using Dremio as a data lake query engine tool. We were creating our PDSs and VDSs on top of our S3 buckets, and our data lake and the data-scientist teams were using the data for further processing. We didn't use it for any ETL jobs. We were using it as a data-lake tool..”

Verified user

Database Engineer at a tech services company with 201-500 employees

[Read full review](#) 

“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. Dermio is as good as any other Presto engine..”

Victor A.

Senior Software Engineer Technical Lead at Apple

[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](#) 

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 ↗](#)

“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..”

Mike Walker

Founder/ CEO at Morphogen Data Inc.

[Read full review ↗](#)

Other Advice

“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](#) 

“We are currently evaluating Dremio against other similar products. But at first glance, I would recommend using Dremio.

Considering my limited access and experience over these three months, I would rate Dremio around a seven out of ten..”

Sachin Shukre

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

[Read full review](#) 

“I rate Dermio 10 out of 10. Data lake 2.0 or 3.0 is gradually replacing data warehouses, but it's not quite there yet. We can't get rid of data warehouses, but it's heading in that direction. Dremio is a phenomenal place to start with that mindset shift and architecture shift. Despite the shortcomings, Dremio has made better progress in moving the industry forward than anybody else out there. .”

MikeWalker

Founder/ CEO at Morphogen Data Inc.

[Read full review](#) 

“My advice to others is if they are creating a data lake for a customer, Dremio would be useful for a data engineering team. If they're willing to create a data lake and they wanted to use it with the cloud-agnostic tool, then it is a good choice. If this solution meets their requirement they should try it out.

I rate Dremio an eight out of ten..”

Verified user

Database Engineer at a tech services company with 201-500 employees

[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](#) 

Top Industries

by visitors reading reviews

Financial Services Firm

30%

Computer Software Company

9%

Manufacturing Company

6%

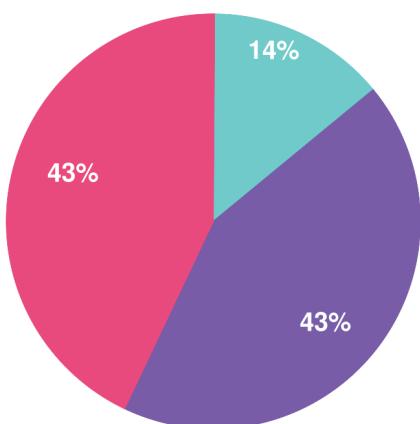
Healthcare Company

5%

Company Size

by reviewers

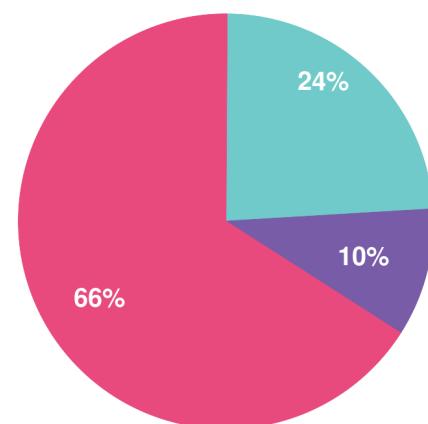
by visitors reading reviews



Large Enterprise

Midsize Enterprise

Small Business



About this buyer's guide

Thanks for downloading this PeerSpot report.

The summaries, overviews and recaps in this report are all based on real user feedback and reviews collected by PeerSpot's team. Every reviewer on PeerSpot has been authenticated with our triple authentication process. This is done to ensure that every review provided is an unbiased review from a real user.

Get a custom version of this report... Personalized for you!

Please note that this is a generic report based on reviews and opinions from the collective PeerSpot community. We offer a [customized report](#) of solutions recommended for you based on:

- Your industry
- Company size
- Which solutions you're already considering

The customized report will include recommendations for you based on what other people like you are using and researching.

Answer a few questions in our short wizard to get your customized report.

[Get your personalized report here](#)

About PeerSpot

PeerSpot is the leading review site for software running on AWS and other platforms. We created PeerSpot to provide a trusted platform to share information about software, applications, and services. Since 2012, over 22 million people have used PeerSpot to choose the right software for their business.

PeerSpot helps tech professionals by providing:

- A list of products recommended by real users
- In-depth reviews, including pros and cons
- Specific information to help you choose the best vendor for your needs

Use PeerSpot to:

- Read and post reviews of products
- Access over 30,000 buyer's guides and comparison reports
- Request or share information about functionality, quality, and pricing

Join PeerSpot to connect with peers to help you:

- Get immediate answers to questions
- Validate vendor claims
- Exchange tips for getting the best deals with vendor

Visit PeerSpot: www.peerspot.com

PeerSpot

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

reports@peerspot.com

+1 646.328.1944