



Yellowbrick 

## Meet the Yellowbrick Cloud Data Warehouse

Unlock value with  
deeper, more meaningful  
insights from your data.



Yellowbrick Cloud Data Warehouse is a Massive Parallel Processing (MPP) analytics database that helps you transform your business by achieving insights in real-time. With its cloud-agnostic architecture, Yellowbrick provides a flexible, future-proof platform for faster innovation.

## Simple, scalable, affordable: Meet Yellowbrick



### Grows with your needs

We offer a unique-to-market cloud-agnostic architecture that will accommodate on-premises, public cloud, private cloud, or a hybrid model of data storage and analytics. We also employ a unique cloud data warehouse that can easily grow and/or contract according to your needs.



### Real-time analytics at any scale

No matter the scale and speed of data analytics you require, Yellowbrick can do it 10 to 100X faster than other data warehouses. And we can do it all while managing concurrent users and ad hoc and mixed workloads, providing you real-time data in, well, real-time.



### Ease of use

We offer fast and easy migrations. Built on standards like PostgreSQL, Yellowbrick's design framework plays well with all your tools, data sources, and infrastructure. That means Yellowbrick's platform is both flexible and futureproof, designed to accommodate the innovations of today and tomorrow.



### Cost management

Yellowbrick provides predictable, subscription-based pricing that removes the risk of cost overruns associated with other cloud data warehouses.

# Any volume of data you need in real time

The last few years have seen an exponential increase in data for many organizations. In fact, you may be searching for a new data warehouse solution precisely because the legacy option you have can't support ad hoc and mixed-workload queries for hundreds (or even thousands) of concurrent users, never mind scaling for future growth. Data has become valuable, but it has no value if you can't access and analyze it quickly, efficiently, and in real-time.

Among the pain points organizations face when it comes to data analytics are the following:

- Queries taking too long.
- Lack of access to real-time analysis of real-time data.
- Inability to run ad hoc queries in timely fashion.
- Performance lags due to large numbers of concurrent users.
- Inability to scale current data analytics solutions.
- Security, performance, and data migration issues resulting from data storage on public clouds.

## What's the solution? A cloud data warehouse like Yellowbrick.

Yellowbrick allows organizations like yours to run analytics in multiple public and private clouds or on-premises simultaneously, offering analytics where and when you need them. That means you get insights in real-time and can make immediate decisions without waiting on slow business analytics performance.

With unmatched performance on petabytes of data, Yellowbrick can support thousands of concurrent users performing complex ad hoc queries and running sophisticated reports. And with its unique-to-the-market cloud architecture, Yellowbrick can readily manage large, difficult-to-migrate datasets where they reside, whether that's on-premises or in the cloud.

Our customers have experienced workloads running tens or hundreds of times faster per node compared to cloud-only or legacy data warehouses they employed previously. At Yellowbrick, we've optimized data processing not only in our own database software but also at the operating system level. We've introduced new networking protocols to efficiently exchange data between cloud nodes in our massively parallel clusters. You get an MPP data warehouse that can accomplish much more at the same price point as a latency-prone, legacy.

The bottom line? Yellowbrick provides your company with a data warehouse with an analytical engine that can process any volume of data in real-time.

# Cloud elasticity and flexibility...

## anytime, anywhere

You may find your organization can't take full advantage of powerful business intelligence and analytics tools because your legacy data warehouse is too inflexible, too expensive to scale, and too difficult to manage. And that hampers your capacity to access real-time insights to make real-time decisions.

Yellowbrick offers cloud deployment flexibility with its innovative data warehousing platform for multi-cloud, hybrid, or on-premise deployments. And you can control all of it via our simple, unified control panel Yellowbrick Manager. Yellowbrick overcomes the performance and scaling limitations of most legacy cloud data warehouses while also keeping costs under control.

How do we do it? At Yellowbrick, we employ a unique cloud data warehouse that helps you easily grow and/or contract your data warehouse as needed. Plus, we offer full deployment flexibility: You can deploy our data warehouse solution in any public or private cloud, on-premises, or a hybrid of all three. This distributed cloud model dramatically reduces (and even eliminates) latency, provides greater data sovereignty and security, addresses data gravity

requirements, and provides uniformity among infrastructure and services. This model also provides for analytics applications at remote locations that may not have internet connectivity (or, thus, access to the public cloud).

Yellowbrick gives you ultimate flexibility in putting analytics where you need them when you need them.

Many legacy and cloud-only data warehouse solutions often create as many problems as they seem to solve. While plenty of data warehouses can adequately run on the public cloud, the real problem manifests itself when you need a hybrid or multi-cloud deployment.

No matter what cloud data warehouse you're using right now, you are likely facing ongoing cost overruns. That's because whether you're using AWS Redshift, Snowflake, or another of the most common options on the market, you're operating with inefficient software combined with a consumption-based pricing model.

## End cost overruns... forever

While consumption-based pricing seems like a win—after all, you only pay for what you use, right?—unless you have intimate knowledge of department-level and user-level behavior, chances are you've found storage is cheap but computing costs are unpredictable and quickly skyrocket.

With Snowflake, for example, you purchase credits up front that you'll burn through in no time because you can't predict how long it will take to run queries in a mixed workload environment with lots of concurrent users. Pretty soon, you'll find your data scientists, engineers, and business analysts restricting their own innovation capabilities because they'll wonder if your organization can really afford to perform some new analysis they need.

So how do you manage costs? With Yellowbrick's subscription model where you can accurately budget your data warehouse costs and use the data without any fear of cost overruns.

### How do we do it?

- You pay for your own storage and computing costs. At Yellowbrick, we don't lease underlying infrastructure from another cloud service provider the way consumption-based subscription models do, for example. Instead, you use your own public cloud provider bulk purchase agreements or credits without an intermediary.
- You pay a flat subscription over a one- or three-year term for the amount of computing you need. This allows you to use Yellowbrick as much as you like, 24/7, 365 days a year without any quotas, limits, or cost overruns.
- You'll enjoy easy and affordable scaling. Need to grow bigger? Yellowbrick can add nodes non-disruptively and will charge you a prorated amount for any extra nodes you need. Our software is so efficient, you will only need a fraction of the compute that other cloud data warehouses will require to match your varying workloads, and so elastic scaling is much more affordable with Yellowbrick.

**With Yellowbrick, you'll see your overall cloud spend drop dramatically in no time.**

# Simple migration, easy adoption

- Complex scripting requirements and lots of specialized personnel to manage them.
- Slow, multi-step, and high-risk migrations.
- Complicated “lifts and shifts,” especially for organizations with on-premises data warehouses that are forced to adjust to a cloud-only solution.
- Inability to provide out-of-the-box performance without extensive fine-tuning.

At Yellowbrick, we eliminate these headaches first and foremost because we’re not limited to public cloud-based architecture like other leading cloud data warehouses. Is your data stored in a private data center or across some kind of on-premises and cloud model hybrid? It doesn’t matter. We’re cloud-agnostic, so we won’t inhibit your organization’s goals with cumbersome and costly migrations. Whether you run in AWS, Azure, Google Cloud; on a private cloud; or on-premises, we can keep your data where it is and let you manage it all via a single control panel.

Yellowbrick provides proof of concept (POC) and performance and operational testing to make sure our cloud data warehouse is right for your organization. Collaborating closely with your team, we’ll work to understand the requirements and scope of a POC, whether in the cloud or in

your on-premises data center. During that POC, we will set up the test instance, migrate your data, and integrate all applications. Set up often takes no more than a day, and we can usually complete performance and operational testing in a week, including the following:

- Loading data from databases, HDFS, or object stores.
- Moving your ETL work.
- Running your SQL queries.
- Migrating stored procedures, embedded SQL, and business logic.

Because Yellowbrick is built on standards like PostgreSQL it will work with all your tools, data sources, and infrastructure not only today but well into the future. Yellowbrick works out of the box with most common industry tools that speak ANSI SQL (like Tableau, MicroStrategy, SAS, and Microsoft BI) as well as with Python and R programming languages.

Whether you currently use AWS, Snowflake, Teradata, Big Query or any of the other leading cloud data warehouses, we have the tools and partners available to help you easily transition and manage cutovers. This means fast, easy migrations and simple user adoption and deployment.

Many organizations are reluctant to abandon their legacy because they know migrations can be complex and costly and prompt user training and adoption hard to achieve. Most leading solutions make setup and migrations difficult, to say the least with a host of potential challenges...

# A different kind of data warehouse

In fact, we have countless examples of new Yellowbrick clients who have been able to run queries they found impossible to perform with their former cloud data warehouse. And many have saved millions of dollars in the process.

Yellowbrick's cloud-agnostic architecture provides your organization with a solution like no other with a range of critical benefits:

- Ownership of your own data and infrastructure with a true single tenant solution.
- Access to real-time data in real-time.
- Multi-cloud, hybrid or on-premise deployments.
- Scaling capacity and flexibility with no fear of cost overruns.
- Simple and affordable subscription-based pricing.
- Fast and easy migrations and deployment with instant out-of-the-box performance.

Yellowbrick's analytics database combines everything you love about on-premises and cloud data warehouses, like industry-leading speed and concurrency for batch, real-time, ad hoc, and mixed workloads, with agility, and a predictable subscription-based pricing model.

**Intrigued?** Find your biggest, most complex data set and most important business challenge and invite us in for a Yellowbrick demo. We'll show you how high-speed data analytics can redefine your business opportunities.

Book a Demo with us:  
Phone 877.492.3282  
[info@yellowbrick.com](mailto:info@yellowbrick.com)  
[sales@yellowbrick.com](mailto:sales@yellowbrick.com)



Yellowbrick is a different type of data warehouse. We employ an architecture that provides complete deployment flexibility with an intense focus on speed, scale, ease of management, and compatibility.