

Reduce Cloud Data Warehouse Costs with Yellowbrick Data Warehouse

A recent survey by Gartner* showed that, “77% of enterprises have been “surprised” by incidents in which costs suddenly spiked.”

“More than a third of those respondents were unable to detect a surge for days, incurring extra costs of as much as \$100,000 and driving their cloud budgets by 10% or more.”

*Source: 3 Techniques for Building Cost Resilience Into Your Cloud, by David Wright, December 9, 2021 - ID G00750749


Today cloud data warehouses such as Snowflake and Redshift have become crucial for modern enterprise analytics and applications due to their instant accessibility, scalability, and elasticity. Snowflake, for example, pioneered the “easy button” approach to data warehousing in the cloud, which promotes simple consumption versus cost stability.

As businesses' needs grow, with ever-growing demands, users, data volumes, and applications, analytics costs quickly and unexpectedly spiral out of control, forcing other business priorities to be put on hold. The Yellowbrick Data Warehouse overcomes some of these challenges with a unique architecture designed to maximize efficiency and simplify management.

Yellowbrick delivers everything you would expect from a modern, elastic, cloud, SQL-based data warehouse. It combines SaaS cloud simplicity with performance perfected through years of delivering the highest ROI to customers on-prem. Optimized for AWS, GCP, and Azure compute instances, organizations can achieve dramatic performance gain, lower their compute and database costs, and use a lower cloud footprint to meet their sustainability goals.

Yellowbrick's transparent, flat subscription model and ability to effectively manage mixed workload priorities means businesses can offload complex, always-on workloads from their existing platform without disrupting their cloud strategy.

Yellowbrick's strong customer commitment has led to an NPS score of 82 (Sept. 2022), higher than any other cloud data warehouse provider. This means you get the most out of your Yellowbrick investments from a partner deeply invested in your success, not just in driving consumption.



Yellowbrick Data Warehouse

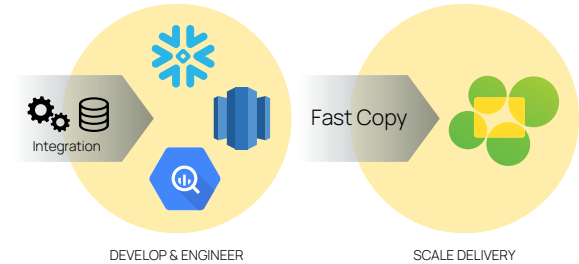
Extreme Efficiency and Performance

Unlike other analytics databases, high-performance real-time streaming data is fully supported and not an afterthought. Many Yellowbrick customers build real-time customer-facing analytics applications. Our deployment model and open instrumentation give service providers the backend they need to build smart multi-tenant analytics applications to monetize data. Our multi-cloud and hybrid approach means you can deliver those capabilities wherever your customers need them, and our unique pricing model keeps the cost to serve low.

Cure Analytics Anxiety

If you are already experiencing cost overruns and cost pressures causing you "analytics anxiety," then deploying Yellowbrick alongside your existing data warehouse is one way to boost performance while lowering overall solution costs. For example, by using Snowflake for its excellent data-sharing capabilities, data engineering, low concurrency, and ad hoc applications, and Yellowbrick for data-intensive, high-concurrency, time-critical workloads, organizations can safeguard their ROI and eliminate their unexpected cloud cost.

Yellowbrick's efficient processing engine and unique, predictable pricing model combine to solve cost challenges and soothe analytics anxiety.



Optimize your data warehouse investment with Yellowbrick

Predictable Pricing Model

Yellowbrick delivers the same deployment simplicity as any other cloud data warehouse. However, with Yellowbrick, you don't need to worry about buying opaque credits or if those credits will run out or get canceled if not renewed. With Yellowbrick, choose the pricing plan you want: on-demand or subscription. For subscription, decide what size data warehouse you want (vCPUs) and pay a flat subscription over a 1- or 3-year term.

This enables you to run applications 24x7 and 365 days a year without worrying about cost overruns.



Yellowbrick.com
Email: info@yellowbrick.com
North America 877.492.3282
Europe +1.650.687.0896

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

Copyright © 2022, Yellowbrick Data, Inc.
All rights reserved.

This document is provided for information purposes only and not warranted to be error free. It is subject to change without notice.

Yellowbrick is a registered trademark of Yellowbrick Data, Inc. Other names may be trademarks of their respective owners.

Extreme Efficiency and Performance

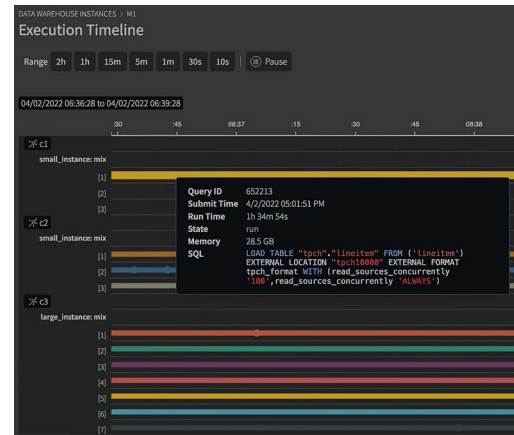
Powered by Direct Data Path technology, Yellowbrick is architected from the ground up for modern applications. Direct Data Path shortens and optimizes data flow to the CPU, delivering extreme efficiencies in query processing from standard cloud configurations.

The result: queries that run faster using far less compute to achieve the same performance as competing solutions. Unlike other cloud data warehouse solutions that simply rely on expensive auto-scaling for performance, Yellowbrick works smarter. This level of efficiency allows Yellowbrick to seamlessly support highly concurrent, ad-hoc queries by thousands of users across complex schemas and fast-changing data.

Yellowbrick is optimized to leverage the performance acceleration capabilities available on AWS, GCP, and Azure compute instances.

Extremely Smart

Yellowbrick's advanced workload manager allocates resources to queries based on policies (out-of-box or custom) to deliver consistent, repeatable, and predictable performance for queries, streaming data, and data ingest. The workload manager ensures critical processes are not starved of resources and can penalty box queries that take longer than their resource policy allows. As a result, you can meet your SLAs while executing complex workloads without automatically adding scale and cost.



Yellowbrick Manager: Visualize workload manager task assignments

Summary

Yellowbrick Data Warehouse provides a low-cost solution to complex, large-scale, mission-critical analytical problems. Optimized for AWS, GCP, and Azure compute instances, organizations can achieve dramatic performance gains, lower their compute and database costs, and use a smaller cloud footprint to meet their sustainability goals.

