Yellowbrick

Stacking up versus Teradata

Yellowbrick Data Warehouse Overview

The Yellowbrick Data Warehouse delivers a modern data warehouse experience in your own cloud account as well as on-premises. We support all the benefits of truly cloud-native data warehousing – separate storage and compute and built-in elasticity, managed through SQL – without your data ever needing to leave your VPC or data center.

Yellowbrick customers are the world's largest global enterprises, running ad-hoc analytics, and operational data warehouses supporting business-critical workloads with high concurrency. Wherever you run Yellowbrick, you'll find rich workload management, real-time data ingest, the ability to load and query data together on the same instance, high availability, and replication across clouds and on-premises for disaster recovery. You can say "goodbye" to the instability of Hadoop and "hello" to a stable, reliable, and trustworthy data warehouse.

Pricing with Yellowbrick is open, simple, and predictable. We've saved many customers millions of dollars per year in legacy and cloud spend. You use and pay for your storage and compute, making use of your cloud credits without paying

smaller companies to mark up your infrastructure costs. You can be comfortable meeting your security, regulatory, and compliance requirements.

If you're considering a new migration to the cloud, Yellowbrick avoids lock-in to any particular cloud vendor or database. We use open-source PostgreSQL as our SQL syntax of choice and are backward compatible with on-premises ETL tools like Informatica PowerCenter and CDC tools, Oracle GoldenGate or Qlik (Attunity) Replicate. We don't do lock-in. We don't want to be the next Oracle (nor did our founders come from there) and we use open standards to ensure the portability of your data and workload.

We're proud to have the best performance in the industry, at the lowest possible cost. No one runs data warehousing workloads faster than Yellowbrick.

Keep reading to see how we stack up versus our competitors in the market.

Yellowbrick

Teradata

We've done away with the things that annoy you the most about Teradata: Fork-lift upgrades, massive on-premises appliances, complex database administration, and – most importantly – the sky-high prices... while supporting the things you've grown to love the most about Teradata: Support for complex, mixed workloads running at high concurrency, predictable response times, a broad portfolio of ecosystem integrations, sophisticated rules-based workload management, advanced SQL support, and the SLAs needed for operational, business-critical workloads.

From there, we've plugged key feature gaps like built-in, asynchronous replication for disaster recovery and support for real-time data streams, and we offer superior query performance at a much lower cost. We wrap this in a modern, truly elastic, cloud-native data warehouse that runs same the in the cloud and on-premises.

Migration is far more straightforward than it used to be. We have tools that automatically migrate the vast majority of stored procedures, BTEQ scripts and ETL, and services that can help with the remainder. We can even survey your existing system up-front to tell you how long migration will take.

A top 10 financial services company compared a 15-node, 6U Yellowbrick system against a 20-rack Teradata 6800. Today, every new transaction across the company's lines of business is loaded into Yellowbrick, instantly queryable for thousands of daily users.

	Yellowbrick	Teradata
Low cost	✓	✗
Autonomous without indexing/tuning	✓	✗
Scale without downtime	✓	✗
Real-time streaming data ingest	✓	✗
High concurrency, predictable latency for operational workloads	✓	✓
Workload management, ad-hoc query	✓	✓
Cloud-native Kubernetes architecture – run in your private cloud as well as on-premises	✓	✗
Built-in asynchronous replication for disaster recovery	✓	✗
Support for semi-structured & geospatial data	Roadmap	✓
Integrated machine learning	Roadmap	✓
Hardware instance size per 1PB warehouse	24 U	16 racks