













Comparing Yellowbrick Cloud Data Warehouse with Teradata

Yellowbrick Cloud Data Warehouse is a DWaaS that runs in your own VPC and 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 your data center.

We’ve done away with the things that annoy you the most about Teradata—forklift upgrades, massive on-premises appliances, complex database administration, support outsourced to IBM, and, most important, 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 much lower cost. We wrap this in a modern, truly elastic data warehouse based on Kubernetes that runs the same in the cloud and on-premises.

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

Category	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		
Kubernetes architecture - run in your private cloud		—
Built-in asynchronous replication for disaster recovery		—
Support for semi-structured & geospatial data	Roadmap	
Integrated machine learning	Roadmap	
Hardware instance size per 1PB warehouse	24U	16 racks