



Yellowbrick Stacking up versus Google BigQuery





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.



Google BigQuery

The Yellowbrick Data Warehouse combines the best features of Google BigQuery with a true, enterprise-grade database.

At Yellowbrick, we've built a sophisticated cost-based query optimizer to enable complex queries across highly variable schemas. We've coupled that with the world's most efficient database execution engine that thrives on running workloads consisting of highly concurrent, complex queries that need to be completed with predictable latency. Google forces you to pick between

“batch” and “interactive” mode – with different tradeoffs and pathetic concurrency limits. At Yellowbrick, we support advanced workload management for thousands of concurrent users to make sure that business-critical SLAs can be met. We even have asynchronous replication – between clouds or between on-premises installations and the cloud – built into the database.

We're proud to have the best performance in the industry, at the lowest possible cost with predictable pricing that doesn't take a Ph.D. to understand or leave you with a hole in the wallet due to a runaway query. No one runs data warehousing workloads faster than Yellowbrick.

	Yellowbrick	BigQuery
Predictable transparent pricing, low cost	✓	✓
Advanced cost-based query optimizer for complex schemas	✓	✗
Cloud data warehouse in your VPC on Azure, AWS, and GCP	✓	✗
Runs on-premises and multi-cloud	✓	✗
Kubernetes-based for private cloud	✓	✗
High concurrency, predictable latency for operational workloads	✓	✗
Workload management	✓	✗
Built-in asynchronous replication for disaster recovery	✓	✗
Compatibility with on-premises ETL (Informatica PowerCenter, etc.)	✓	✗
PostgreSQL-compatible syntax	✓	✗
Elastic with separate storage/compute	✓	✓
Full control over network path and firewalls	✓	✗
Integrated machine learning	Roadmap	✓
Support for semi-structured & geospatial data	Roadmap	✓

