



# Yellowbrick Stacking up versus Amazon Redshift





## 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.



# Amazon Redshift

If you've been looking at or using Amazon Redshift, you must be intrigued by the benefits of true elasticity and cross-cloud portability of competing solutions, yet worried about more vendor lock-in, unpredictable performance, and run-away cost. The Yellowbrick Data Warehouse takes the best of both worlds. Our data warehouse is autonomous, requiring little to no DBA involvement, inherently multi-cloud (and on-premises), has simple transparent pricing, and has built-in asynchronous replication for disaster recovery.

We do all of this without the quirks and complexities that come with a 15-year-old code base like Redshift: We support higher concurrency and mixed workloads without having to worry about things like "vacuuming" or if your cluster can scale without downtime. You can take CDC or Kafka data and stream it into Yellowbrick in real-time, with great performance, and our SLAs are designed for operational, business-critical workloads.

	Yellowbrick	Redshift
PostgreSQL-compatible syntax	✓	✓
Low cost	✓	✗
Runs in GCP, AWS, and Azure	✓	✗
Runs on-premises	✓	✗
Autonomous without indexing/tuning	✓	✗
Built-in asynchronous replication for disaster recovery	✓	✗
Real-time streaming data ingest	✓	✗
Scan acceleration support	✓	✓
Support for semi-structured & geospatial data	Roadmap	✓
Compatibility with on-premises ETL (Informatica PowerCenter, etc.)	✓	✗
Automated migration from legacy technology	✓	✓

