



# Yellowbrick

## Stacking up versus Snowflake





## 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.





# Snowflake

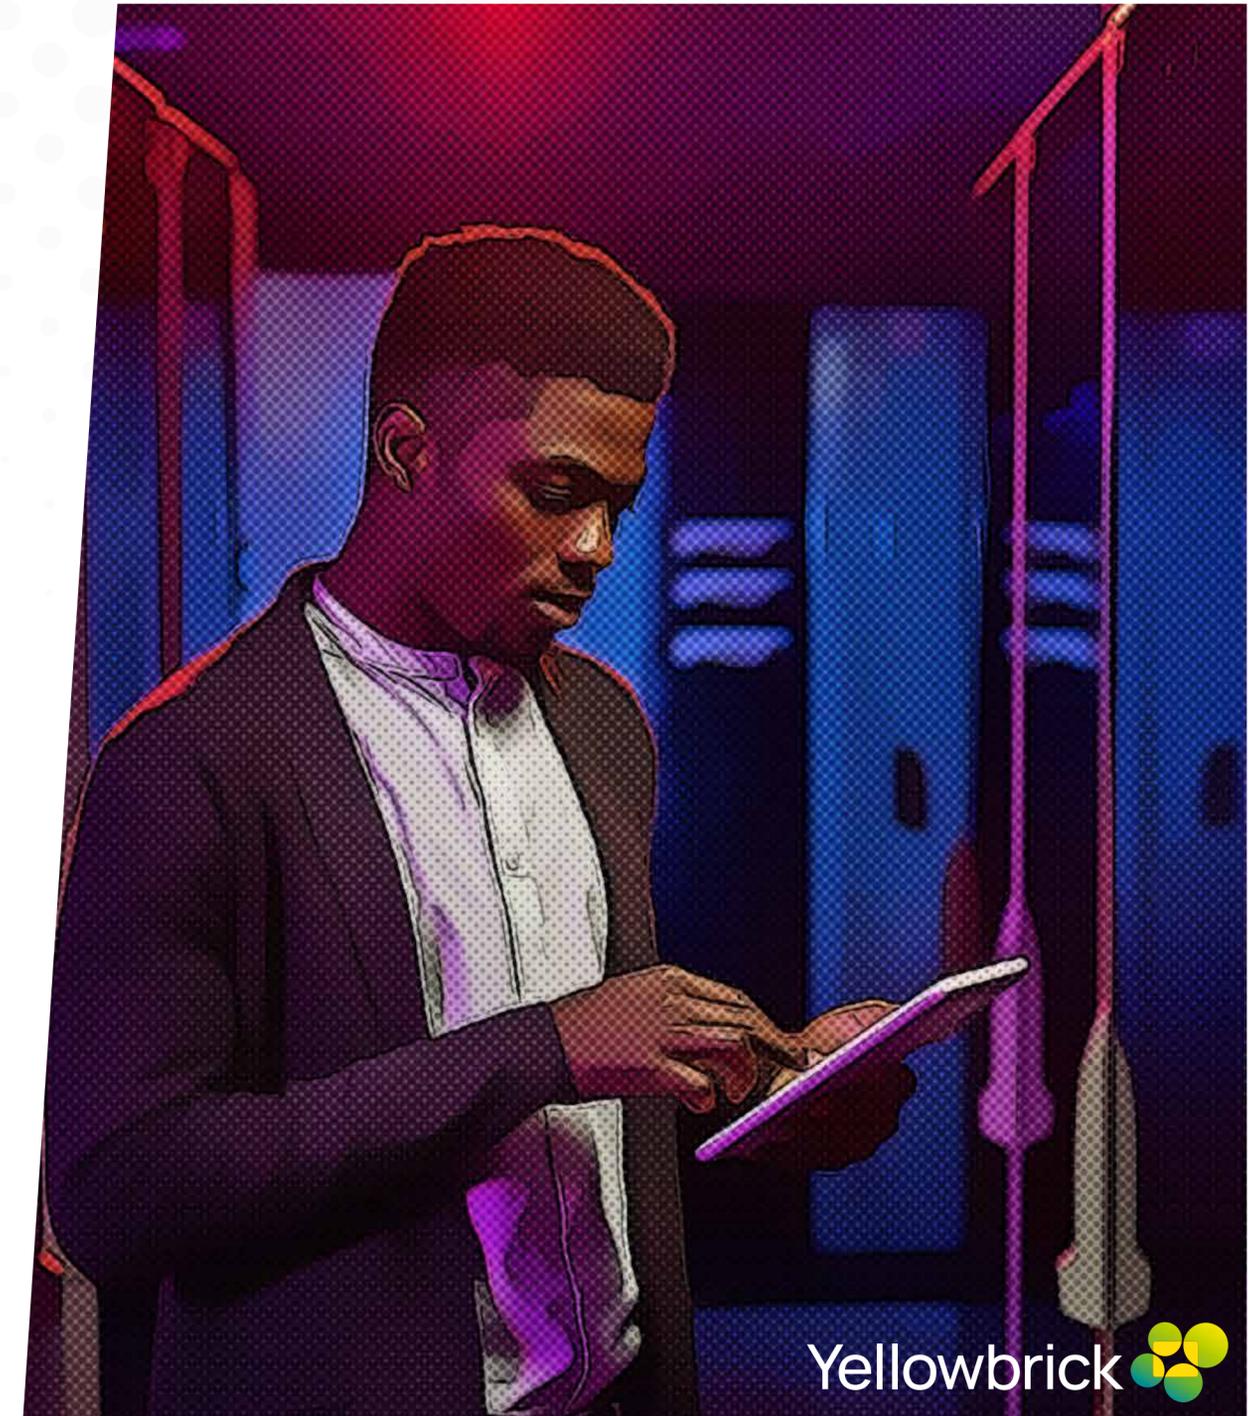
Yellowbrick runs both on-premises as well as in your own cloud account. Snowflake is a cloud-only offering. For many enterprises, the sensitivity of their data and workloads, as well as factors such as data gravity, means that not all data warehousing operations can be moved fully to the cloud.

Our pricing model is predictable – you won't need to worry about buying proprietary credits, or if they will get canceled on you if you don't renew. We don't have Snowflake's low query concurrency limits that force you to run even more clusters to burn even more credits. With Yellowbrick, you simply decide what size data warehouse you want (vCPUs) and pay a flat subscription over a 1- or 3-year term. You can consume as much of Yellowbrick as you like, 24x7 and 365 days of the year, without worrying about cost overruns.

Snowflake scales in large expensive increments. With Yellowbrick you add one node at a time and are charged by the second for the extra vCPUs – no huge step changes in cost. With our advanced workload management, you'll end up needing to scale much less frequently and needing a fraction of the infrastructure you'd need to get the equivalent performance for mixed workloads in Snowflake, greatly reducing spend.

With Snowflake you lose ownership of your data and the infrastructure it runs on. Yellowbrick runs in your own cloud account and network. You benefit from your enterprise agreements with cloud providers. And still, we've made the process of operating Yellowbrick easy – it's data warehousing as a service.

Nobody runs data warehousing workloads faster than Yellowbrick, and we've saved some Snowflake customers millions of dollars per year.



# Snowflake

A digital marketing customer conducted query performance tests using 12 representative SQL workloads against its own data. The data consisted of 3.5 trillion records distributed across 82 tables, with a volume of approximately 500 TB. The query tests compared a 32-node 2XL Snowflake cluster with a 30-node Yellowbrick cluster. Yellowbrick was 6x faster on average across all 12 queries and 147x faster on one of the longest running queries.

	Yellowbrick	Snowflake
Predictable, low cost	✓	✗
Consumption pricing with dollars, not credits	✓	✗
Per-second billing	✓	✓
Runs and secured in your own VPC/cloud account	✓	✗
On-premises deployments	✓	✗
Disaster recovery & protection from global outages	✓	✗
Elastic with separate storage/compute	✓	✓
Kubernetes-based for private cloud	✓	✗
High concurrency, predictable latency for operational workloads	✓	✗
Workload management	✓	✗
Real-time streaming data ingest at no additional cost	✓	✗
Compatibility with on-premises ETL (Informatica PowerCenter, etc.)	✓	✗
Full control over network path and firewalls	✓	✗
Granular scale by single node	✓	✗
Integrated machine learning	Roadmap	✓
Data sharing	Roadmap	✓
Support for semi-structured & geospatial data	Roadmap	✓

