



Yellowbrick

Stacking up versus Netezza





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.



If you’ve used the IBM Netezza platform, IBM will try to sell you “Cloud Pak for Data with Netezza Performance Server.” IBM has succeeded in positioning this as a new, modern data warehouse offering; but behind the charade is the ancient, buggy Netezza database engine with the same old limitations: Poor concurrency; poor write performance; largely missing workload management; weak business continuity; a support team that can’t debug or fix core issues; performance that’s weak compared to modern cloud offerings; no working replication or disaster recovery strategy; and most importantly, no real roadmap at all.

IBM will position this “new” offering as working in the cloud, but it’s just a “lift and shift, bundle in

containers” of legacy technology built on Postgres 7 from the year 1999. Truly cloud-native databases have separate storage and compute, storing persistent data on object stores like S3. IBM doesn’t. Truly cloud-native databases have SQL-managed elasticity and modern, web-based query tools. IBM doesn’t. Truly cloud-native databases authenticate with OAuth2. IBM doesn’t. The development teams of truly cloud-native databases believe that now is the most exciting era in database architecture in decades – they don’t put a relic in a container and re-spin it as a reincarnation of a dead brand.

If you’re going to invest in a new data warehouse, you want Netezza compatibility, but with a truly cloud-native architecture, elasticity, a modern management experience, high concurrency, business continuity, the ecosystem, and ETL integrations you know and love – that runs in any cloud and on-premises. There’s no need to look at IBM anymore: Choose Yellowbrick.

“In our testing of Yellowbrick, we compared the performance of a six-rack TwinFin to the 6U baseline Yellowbrick system. And performance was anywhere from 3 to 50 to 100 times faster.”

RICK MAHUSON Vice President, R&D Business Analytics, TEOCO

	Yellowbrick	Netezza
High software reliability and excellent support	✓	✗
Actively maintained database engine with a roadmap	✓	✗
Cloud data warehouse in your VPC	✓	✗
Elastic with separate storage/compute	✓	✗
Data stored on S3, no backups required	✓	✗
High concurrency, predictable latency for operational workloads	✓	✗
Workload management	✓	✗
Real-time streaming data ingest	✓	✗
High density, small footprint solution	✓	✗
Runs in AWS, your VPC	✓	✓
Runs on-premises	✓	✓
Autonomous without indexing/tuning	✓	✓
Scale without downtime	✓	✗
Built-in asynchronous replication for disaster recovery	✓	✗