The End of An Era

What's Next for Netezza Users?

Modernizing IBM Netezza with Yellowbrick Data Warehouse

Overview

Despite being the foundation of many corporate analytics environments, IBM has failed to invest in the Netezza data warehouse appliance. Instead, they treat it as a "cash cow." The future is bleak for Netezza users. No investment in a technology roadmap and ever-changing guidance about End of Life (EOL) introduce unacceptable risks for organizations dependent on this once-excellent technology. Today, IBM Netezza users live on borrowed time before the next SLA breach or service outage occurs.

Netezza's first customer and Netezza's largest customer have both migrated to Yellowbrick, along with many other long-time Netezza customers including Bread Financial, Catalina Marketing, Nielsen, Symphony RetailAI, Telcel, and the US Navy. Yellowbrick offers a better, faster, and more cost-effective solution with two massive advantages: cloud and a roadmap.

In the cloud, Yellowbrick offers cost-effective full elasticity with separate storage and compute on AWS, Azure, and Google Cloud as well as on-premises. Yellowbrick's database already possesses capabilities that IBM should have been offering to its customers, such as production-grade asynchronous replication, real-time data ingest, and more concurrency, and has an aggressive roadmap around data sharing and machine learning in the cloud as well as on-premises.

The Yellowbrick Cloud Data Warehouse

The Yellowbrick Data Warehouse is an elastic, massively parallel processing (MPP) SQL database designed for the most demanding batch, real-time, ad hoc, and mixed workloads. It can run complex queries at up to petabyte-scale with guaranteed sub-second response times. Yellowbrick is proven to provide business-critical services at many large global enterprises with thousands of concurrent users. The front-end database service, like Netezza, is compatible with PostgreSQL, albeit a much more modern version.

Better

Yellowbrick delivers enterprise data warehouse functionality for business-critical needs on the scale necessary to handle large workloads reliably, offering a step up from Netezza.

Optimized Price/Performance

Yellowbrick's patented Direct Data Accelerator Architecture is an OS bypass technology, enabling in-memory analytics performance at petabyte scale without requiring a typical database buffer cache – leading to more predictable response times and massive cost reductions.

Yellowbrick

Yellowbrick dramatically shortens time-to-insight over Netezza, with some customers reporting up to 100x improvement. Replacing Netezza with Yellowbrick shatters limits on analytics across petabytes of data with few modifications needed.

	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	✓	✗

Predictable and Transparent Pricing

Yellowbrick supports both on-demand and subscription-based pricing. All pricing is based on the consumption of vCPUs for compute; we do not charge for storage since data is persisted on object storage in the customer's own cloud account (or as part of an on-prem hardware subscription from Yellowbrick). On-demand pricing caters to short-term burst needs and is billed monthly in arrears. Subscription pricing is predictable, works across cloud and on-premises, and allows efficient acquisition of capacity when needed. Models can be mixed and matched to meet business objectives. With Yellowbrick pricing, customers avoid paying for expensive systems at idle and can grow incrementally without lift and shift.

Easy Migration

Like Netezza, Yellowbrick on the surface looks like Postgres, massively simplifying migration, and we have fully supported ecosystem partnerships with many of the same ETL, BI, and CDC vendors as IBM. Being largely dialect compatible with Netezza, migration of workloads is straightforward; even the DBA/administration experience is very similar. Yellowbrick partners with KPMG, Capgemini, Accenture, Next Pathway, ZS, Systech, Smart Associates, and Cognizant for other ongoing development and migration work where necessary. Some of these partners offer tools for estimation of migration effort and cost, along with automated migration and testing of the vast majority of database objects and ETL.

Summary

Yellowbrick Data Warehouse provides a low-cost solution to complex, large-scale, mission-critical analytical problems. It runs on any cloud and is charged through a predictable pricing model. Yellowbrick's modern cloud data warehouse is the natural successor to Netezza, delivering high performance, massive concurrency, and real-time ingest running in the cloud or on-premises.

Yellowbrick

Yellowbrick.com
Email: info@yellowbrick.com
North America 877.492.3282
Europe +1.650.687.0896

Copyright © 2022, Yellowbrick Data, Inc
All rights reserved.

This document is provided for information purposes only and not warranted to be error free. It is subject to change without notice.

Yellowbrick is a registered trademark of Yellowbrick Data, Inc. Other names may be trademarks of their respective owners.