



Breaking through Netezza End-of-Life

An executive whitepaper on data warehousing, post-Netezza



Breaking through Netezza End-of-Life

An executive whitepaper on data warehousing, post-Netezza

How to get started with Amazon Redshift	,
5 Key features of Amazon Redshift	,
5 Key benefits of Amazon Redshift	,
Enterprise data warehouse functionality	4
Ad hoc queries on petabyte-sized data sets	4
Sophisticated workload management	4
Real-time data and large data sets	4
Mixed workloads	4
High concurrency	5
Cloud native Amazon Redshift)
PostgreSQL compatibility	5
Support for stored procedures	6
On-premises, off-premises, and cloud deployments	6
Groundbreaking architecture for peak performance in a small footprint	6
Workload management	6
Utilities	6
Fast, simple installation	6
9 Key features of Amazon Redshift	+
6 Key benefits of Amazon Redshift	+
No systems administration	7
Make DBAs lives easier	7
Increase reliability	7
Expand on demand	7
Amazon Redshift vs. Netezza	,



The rise and demise of Netezza

The move by IBM to discontinue support for the Netezza product line has IBM customers facing a hard choice. On one hand, they can undertake a lengthy and complex migration to the IBM PowerPC DB2 mainframe product — one that is far more complex and costly than Netezza users are accustomed to.

On the other hand, Netezza users could consider cloud options, renting a data warehouse solution that might appear to have better economics. However solutions like Redshift from Amazon will cost more when run 24x7 and are unlikely to provide the same type of performance as Netezza, especially for a large number of concurrent users. Another option is Snowflake, but the Snowflake database is largely incompatible with Netezza and cannot be run on-premises. It is also cost-prohibitive for always-on and mission-critical deployments, which is why the solution allows customers to scale down or turn off the data warehouse when not needed.

This has left Netezza and other traditional data warehouse users without options. Until now.

A path forward for Netezza users

Fortunately a new market entrant, Yellowbrick Data, has already helped many customers secure a compatible, high-performance, and capable product for data warehousing in a post-Netezza world. In fact, one Yellowbrick Data customer noted:

“In our testing of Yellowbrick, we compared the performance of a **gjl !fUW`Hk]b:]b** to the **gjl !!` fl \$Vh \][\ L'VUgY]bYMY`ck Vf]W_`grngHya** . And performance was anywhere from **'`tc`) \$`tc` %\$\$`hja Yg`ZUghYf`**.”

- Rick Mahusan, VP R&D Business Analytics, TEOCO [view video](#)

About Yellowbrick Data

Yellowbrick Data was founded in 2014 by experts in database and high-performance hardware to solve the challenges of multi-user analytics for modern enterprise structured data sets. Yellowbrick Data retains the features, reliability, and predictability of mature enterprise data warehouses while offering compelling economics and price points suitable for today's large data volumes.

The Yellowbrick Data Warehouse achieves these objectives, delivering a solution in a compact form factor that is quick to deploy and easy to expand. It provides end-to-end analytics across the hybrid cloud from the data center to the cloud to the edge.

Enterprise data warehouse functionality

Enterprise data warehouses run workloads supporting business-critical processes that require exact answers to queries in a timely, dependable fashion. For financial calculations, it is unacceptable to produce incorrect results, or for patient data, it is unacceptable to recommend the wrong treatments. Data warehouses must be available to answer business-critical queries and compute reports 24/7/365. Hadoop-based solutions have proven incapable of meeting business-critical needs, and single-server databases cannot deliver on the scale necessary to handle large workloads.

Yellowbrick delivers the following mission-critical capabilities:

5 X \ c Wei Yf]Yg`cb`dYhUVntH!g]nYX`XUH`gYhg`

Analytics and data science are now pervasive across the enterprise, which has made it impossible to perform regular index maintenance and define and build cubes ahead of time. Yellowbrick is the only database to offer rapid response to completely ad hoc queries on large datasets without any indexing, vacuuming, or cube building.

Gcd\]gh]WUHX`k cf_`cUX`a UbU] Ya Ybh`

Yellowbrick guarantees that business-critical reports complete on time to meet service level agreements (SLAs) for interactive users without disruption. With other databases, a single poorly-written query can impact other users, hogging resources and stalling business decisions. Cloud-only and open source databases do not have the sophisticated workload management capabilities enterprises require.

FYU!Hja YXUH`UbX``Uf[YXUH`gYhg`

Modern data warehouses need to cost-effectively scale from terabytes to petabytes of data, while still allowing real-time ingest of events so analytics are up-to-the-second. For example, a business request might be to compare what is happening now with what happened at the same time last year. Cloud-only and open source databases cannot accomplish this. Yellowbrick Data is the only data warehouse, on-premises or otherwise, that can analyze petabyte-sized data sets cost-effectively.

A]l YX`k cf_`cUXg`

Modern workloads feed hundreds of thousands of records per second into real-time systems, along with concurrent ETL processes, ad hoc, and batch queries. Yellowbrick runs all of these at the same time, a capability that cloud-only and single-server databases do not have.

Enterprise Data Warehouse

Enterprises with thousands of users need to access the data warehouse with fast response times. Yellowbrick delivers all of the capabilities listed above, all the time, to thousands of users concurrently accessing data. Public cloud-only and single-server databases simply cannot achieve this.

Optimal for Netezza users

The Yellowbrick Data Warehouse is an optimal replacement for Netezza users, meeting or exceeding the Netezza experience in the following areas:

	Similar to Netezza	Better than Netezza
PostgreSQL compatibility	✓	
Support for stored procedures	✓	
On-premises, off-premises, and cloud		✓
Peak performance in a small footprint		✓
Workload management		✓
Utilities		✓
Fast, simple installation		✓
Real-time data ingest		✓

PostgreSQL Compatibility

Yellowbrick Data maintains compatibility with PostgreSQL and adheres to the ANSI SQL standard to provide a quick and easy migration from Netezza. Yellowbrick has also added functions that are similar to Netezza, which makes the transition to Yellowbrick even easier.

Gi ddcfhZcf`ghcfYX`dfcWXi fYg`


Yellowbrick supports PL/pgSQL stored procedures that allow customers to easily port from Netezza to the Yellowbrick Data Warehouse.

Cb!dfYa]gYgžcZ!dfYa]gYgžUbX`Wci X`XYd`cma Yb!g`

Yellowbrick supports deployments for the hybrid cloud with a Yellowbrick cloud solution (coming in 2019). This provides Netezza customers an immediate solution that they can deploy today, while giving them a go-forward hybrid cloud strategy that spans off-premises and on-premises locations.

; fci bXVfYU_]b[`UfW]hYWi fY`Zcf`dYU_`dYfZfa UbW`]b`Uga U``Zc!df]bh

Yellowbrick Data developed a revolutionary new architecture focused on solid-state storage that introduces dramatic miniaturization to the data warehouse industry. Some customers have seen their Netezza solutions shrink in space from six full data center racks to just six rack units (12 inches!) in size, and to do so with better performance.



K cf_`cUX`a UbU] Ya Ybh

Yellowbrick delivers sophisticated workload management capabilities that allow administrators to assign priority to users and workloads, guaranteeing that service level agreements are met.

I H`]H]Yg`

Yellowbrick offers equivalent tools to Netezza. If you like a utility provided by Netezza, there will be an equivalent in Yellowbrick that is even better.

.

Yellowbrick Data Installation

Yellowbrick Data installs in under a couple of hours and can be production ready in under a day.

Easy to test, easy to deploy

With a compact form factor and familiar feel, Yellowbrick Data can be tested and deployed with ease. Yellowbrick Data offers prospective customers the ability to test workloads on a live system, as well as the ability to conduct a proof-of-concept at the customer site before purchase. These initiatives often take just a fraction of the time originally planned, with one customer noting:

“We set up a 3-week POC and we were essentially done at the end of one week.”

- Atul Jain, Chairman and CEO, TEOCO [view video](#)

Slash operational costs

Yellowbrick Database Appliance

Yellowbrick is a self-contained database appliance, just like Netezza, but has an even more compact, simple, and integrated form factor. There's no need to provision storage or maintain file systems; no need for network administration or configuration; no operating systems to install or upgrade; no firmware or BIOS settings; and no Java virtual machines to handle or processes to restart.

Reduced Operational Costs

The database requires no indexes, no cubes, no vacuuming, and no grooming. Many databases supporting multiple customers can run on one system, so your DBA resources can spend more time directly contributing to business results.

Reliability


Spinning hard drives are notoriously unreliable and are the most common service item in any data center. Yellowbrick has an all-flash architecture to maintain system reliability. All on-site service and support is provided by Yellowbrick Data according to the enterprise level SLAs you expect.

Yellowbrick Data is a self-contained database appliance.



9I dUbX'cb'XYa UbX'

Yellowbrick easily and cost-effectively grows alongside customer demand. Rather than having infrastructure teams install new servers, switches and storage, and run complex migration utilities, customers can simply add new blades to the Yellowbrick system in seconds. The system automatically makes the new compute and storage capacity available.



“The performance improvements we saw were from 3x to 10x, running ‘as-is.’ We had six engineers touch the system and all of them found the system very easy to use because there was a lot of commonality with the existing systems that we already had. We’re expanding Yellowbrick out to multiple customers across the globe.”

- Nigel Pratt, SVP Technology, Symphony Retail AI [view video](#)

Make your Netezza migration fast and simple

With the end dates for Netezza products fast approaching in 2019, the time is right to get started today. Visit www.yellowbrick.com or contact us at info@yellowbrick.com to schedule a proof-of-concept or learn more about how we can help your migration fast and simple.



250 Cambridge Avenue, Suite 300, Palo Alto, California 94131, USA, 1.650.687.0896

© Yellowbrick Data, Inc. All rights reserved. Yellowbrick Data and the Yellowbrick logo are trademarks of Yellowbrick Data. All other trademarks are the property of their respective owners. Information is subject to change without notice.