Yellowbrick

COFFING DATA WAREHOUSING

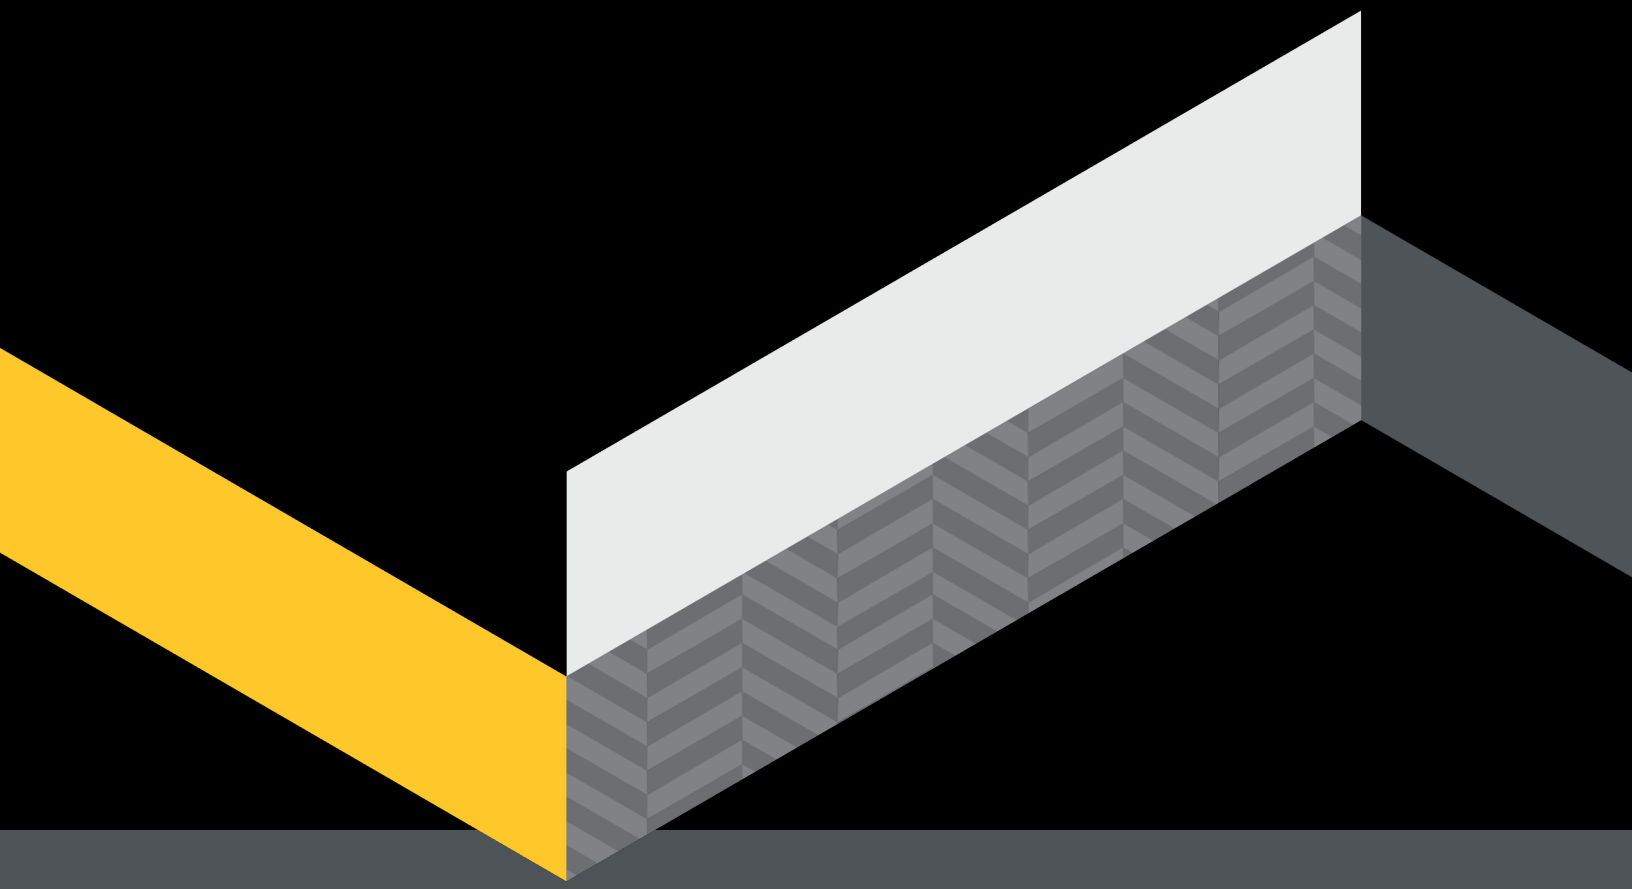

MIGRATION GUIDE

Automated migration from SQL Server to Yellowbrick with Nexus Server

Accelerate migration from SQL Server to Yellowbrick for 100X faster performance at a fraction of the cost

SQL Server is not designed for real-time analytics

Across any organization size or industry, you need access to accurate, lightning-fast insights to remain competitive. That requires a highly optimized, MPP ("scale-out") analytic database that can keep up with increasing data volumes and growing number of users—which SQL Server's "scale-up" architecture was never intended to do.

SQL Server presents several challenges to organizations focused on delivering modern analytics:

Poor scalability: SQL Server cannot scale to support bigger workloads without expensive hardware or dividing applications into multiple data marts, making management a nightmare.

Performance problems: SQL Server performance can be lacking, with reports and queries taking minutes or even hours.

Difficult to budget: A complicated licensing scheme, along with what can be high costs, make SQL Server economics difficult.

Management complexity: SQL Server requires tuning and other optimizations for peak performance, expansion is painful and expensive, and backing up datasets even 1TB in size is problematic.

Weak workload management: With SQL Server, managing workloads for larger use cases is complex and difficult, often leading to performance and reliability problems.

Yellowbrick is purpose-built for modern analytics

Yellowbrick Data Warehouse is an advanced, massively parallel, 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 across numerous nodes, with guaranteed sub-second response times. Yellowbrick was conceived with the goal of optimizing price/performance. It's not uncommon for customers to see their workloads run tens or hundreds of times faster at a fraction of the cost compared to cloud-only or legacy data warehouses.

Yellowbrick continuously implements new hardware and software protocols in an adaptive "cut-through" architecture that ensures the best performance in every environment. Yellowbrick combines these advances with smart thinking about storage formats and indexing, and add on top a modern, standards-based database interface that is familiar to users (PostgreSQL) for ecosystem compatibility.

The result is a modern, quickly provisioned, and easy-to-use data warehouse that delivers the best price/performance economics in the industry with full deployment flexibility (private, public, and hybrid/distributed clouds). Yellowbrick Manager provides a simple, unified control plane to manage all instances, databases, and users.

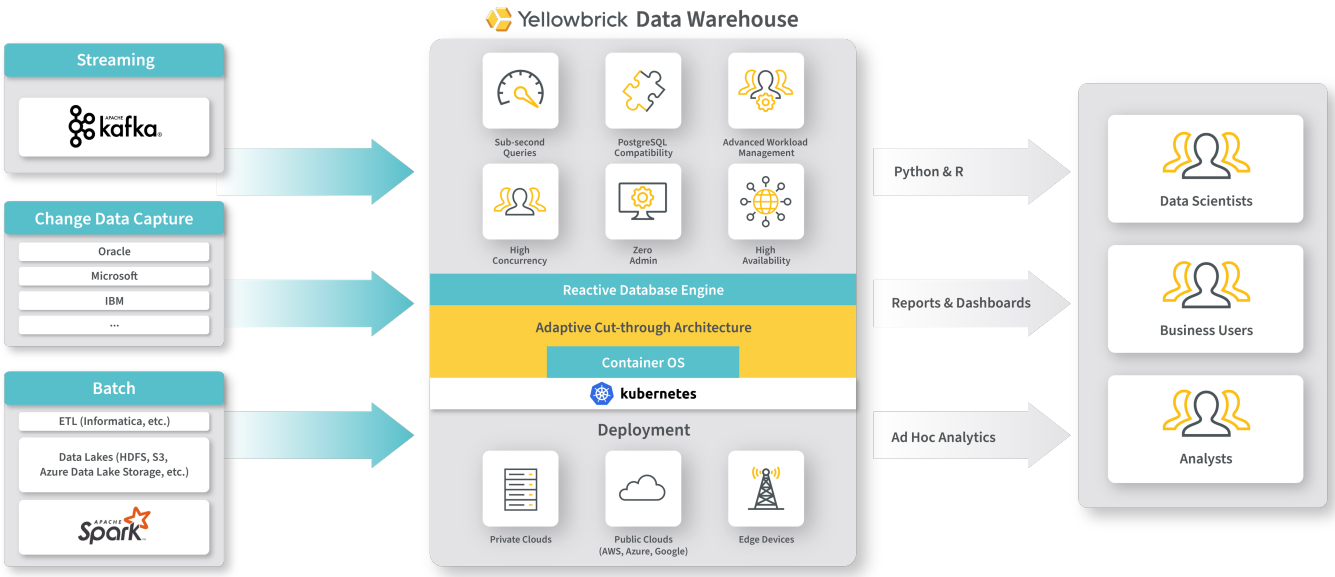

Yellowbrick
Velocity Partner Network

MIGRATION GUIDE

Replace SQL Server with Yellowbrick for 100X performance at a fraction of the cost

Replacing SQL Server databases with Yellowbrick Data Warehouse brings the power of interactive, ad hoc analytics to thousands of users, simplifies operations, and reduces costs.

- **Analyze data 100X faster:** Designed for highly optimized MPP analytics, Yellowbrick enables real-time, ad-hoc queries for thousands of users, far beyond what SQL Server can achieve.
- **Consolidate disparate databases:** Yellowbrick Data Warehouse can scale from single TBs to multiple PBs, offering a way to combine multiple SQL Server databases into a single, easy-to manage instance. Or, easily replicate data from SQL Server to Yellowbrick for Change Data Capture scenarios.
- **Simplify management:** With Yellowbrick, unlike SQL Server, there's no need for manual time-consuming tasks like query tuning or building indexes. PostgreSQL compatibility ensures access to common skill sets.
- **Count on rock-solid reliability:** Yellowbrick's Advanced Workload Management offers granular control of workloads to prioritize concurrent queries across massive amounts of data, bringing added reliability and performance.
- **Grow along with the business:** Yellowbrick instances can be expanded easily without downtime, eliminating any impact on operations.
- **Fit into the Windows world:** Yellowbrick integrates with Microsoft Active Directory, and for cloud deployments, functions as a Private Link Service in Azure. Plus, Power BI and a full range of other enterprise ecosystem tools are supported.

Customer Spotlight

BMW Group Financial Services

BMW Group Financial Services is one of the leading financial services providers in the automotive sector, serving customers worldwide.

Challenge

The customer's SQL Server-based analytics and reporting solution was at the edge of performance and scalability limits for use cases like funding analysis, payments aggregation, and calculating portfolio internal rate of return.

Results with Yellowbrick

BMW Group chose Yellowbrick to replace SQL Server as its analytics and reporting platform.

Test results including:

•152X faster queries on average

•70% reduction in required data storage space

•Successful integrations with Informatica,Tableau, and SAS

"We're confident that Yellowbrick's ability to quickly analyze large amounts of data and offer new insights will help us deliver on our goals."

Ian Smith
CEO, BMW Group Financial Services North America

Fast and simple migration from SQL Server to Yellowbrick

Yellowbrick extends its software capabilities as the only vendor to partner with the Nexus product line from Coffing Data Warehousing. The Nexus Chameleon combines user-friendly desktop software with the Nexus Server to automate conversions and data movement between databases at high speeds. As a result, business users, data scientists, the DBA, and ETL teams can query all systems, migrate data between all systems, and join data across all systems.

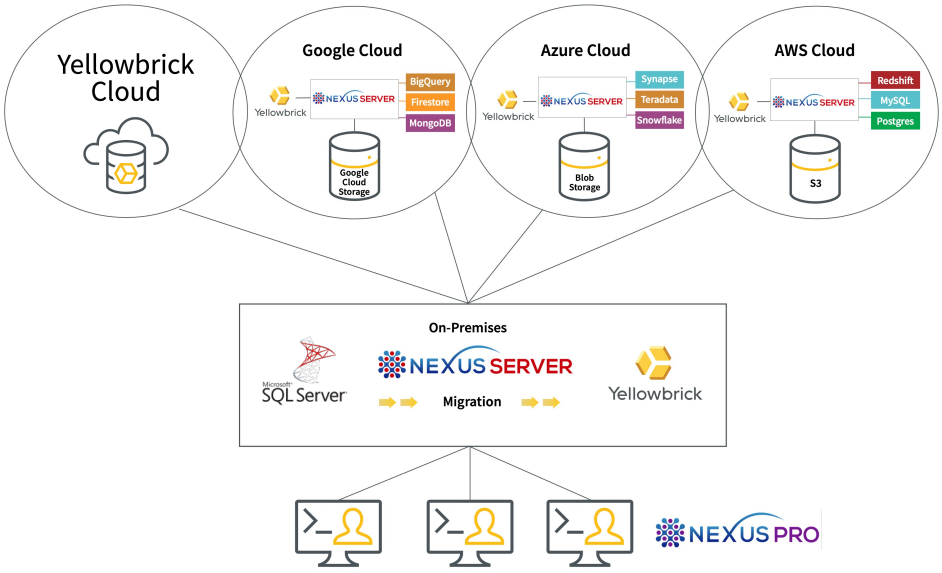

As businesses' data access, migrations, and federation need to grow and expand, workforces become increasingly remote and decentralized. Employee access into the data environment is achievable in several ways, including VPNs and remote desktop solutions, but these strategies are inadequate solutions to the need for efficient data movement. Furthermore, schema conversion, data translation, and ELT processing can be cumbersome and technically challenging even in optimal networking environments.

Nexus Server is a software solution that offers a complete data migration solution between database platforms. A local installation of Nexus Desktop client allows users to manage high volume data movement regardless of their physical proximity to the data. In addition, the Nexus Server converts database schemas and appropriately translates data between database platforms to limit excessive administrative hours.

When you deploy a Yellowbrick data warehouse with multiple SQL Server databases and place a Nexus Server on a high-speed network, you enable a frictionless migration solution. Users running the Nexus Chameleon desktop software can move and join data seamlessly across platforms as if all data is on one centralized data warehouse. Business users no longer need to fill out a ticket so IT can gather data for them because everyone has access to any data, at any time, anywhere.

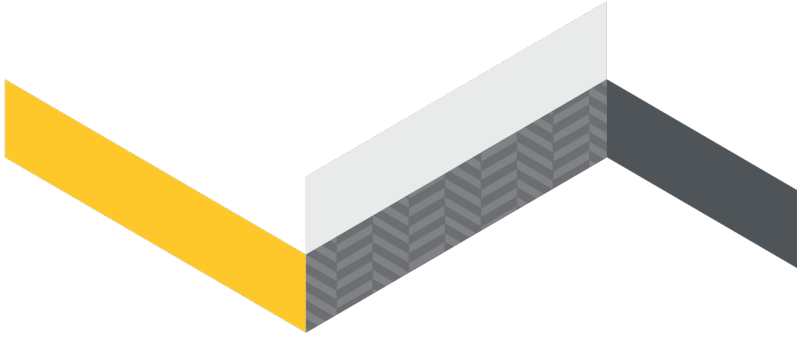

Nexus Server delivers a comprehensive migration solution

One-click migration

Data migration typically requires schema and data type conversion, data translation, and load utility script creation for each table included in the migration.

A project that involves the migration between two database platforms can take months to prepare. In addition, resolving issues and errors during the load process can be time-intensive, and data loss can occur.

Nexus Server assists in migrating SQL Server to any Yellowbrick instance by generating ETL based on the source and target systems and supports custom ELT scripts. Nexus allows users to change the columns and rows they want to move, the target table definition (DDL), indexes, and distribution. It also provides the ability to transform data before or after the migration occurs. However, it can be as simple as choosing the SQL Server database, the tables you want to move, and scheduling it on the Nexus Server.

Flexible execution

The majority of utility-based data migration is processed locally to the user. However, local processing can create issues when businesses rely more on remote workforces and route the data through a throttled VPN or even the public internet.

User-level networking can quickly become saturated by extensive data migration processes competing with other local area network needs, even in on-site migration processes.

Nexus Server allows the execution of data movement in a location as close to the database as possible using networking dedicated to data movement. Working from any site, ETL teams can prepare migrations, business users can transfer data or set up federated queries, and users can schedule every process on a global calendar. In addition, the Nexus Server honors individual user credentials at the database and domain levels to retain user access rights and provides a superhighway for data movement.

Federated queries support

Instead of having a single SQL Server data warehouse, a distributed approach is a good case. Hence, data remains with the line of business, with the domain experts, with the users that created that data. However, with governance and metadata descriptions in place, other departments can use the Nexus Chameleon and Server to access the data, have confidence in the quality of the data, and trust the data.

As businesses' data needs grow and change, data "silos" can naturally develop. Different departments choose database platforms based on their own needs, but this can cause problems when business decisions need information from multiple sources.

The Nexus client software allows users to view all of their database sources in a unified Systems Tree. In addition, users can use the Cache Tree feature to search every connected data source "silo" for metadata such as object names, column names, and data types.

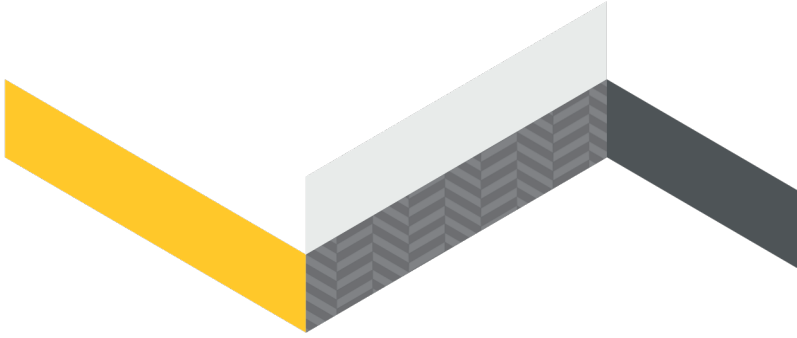

Nexus can execute federated queries, joining tables and views across database platforms deployed on-premises and in the cloud. Users can choose where to process the join, including SQL Server, Yellowbrick, other database platforms, the Nexus Server, or even the user's desktop.

Users can also design reports that pull data from multiple sources in a drag-and-drop interface and share these federated queries with peers. Once a user creates a report, they can schedule it to run or execute it with the press of a button.

Steps to create Utopia for users of data

1. Provision a Yellowbrick instance, a Windows Server, and install Nexus.
2. Use the Nexus Server to migrate specific tables from SQL Server to Yellowbrick.
3. Use the Nexus Chameleon to run queries to compare SQL Server and Yellowbrick.
4. Set up Super Join Builder to automatically join SQL Server and Yellowbrick tables.
5. Use the Nexus Metadata Cache Tree to search object names across all systems.

Nexus Server in action

The image below shows the Nexus Server menu and is the starting point to quickly migrate SQL Server to Yellowbrick.

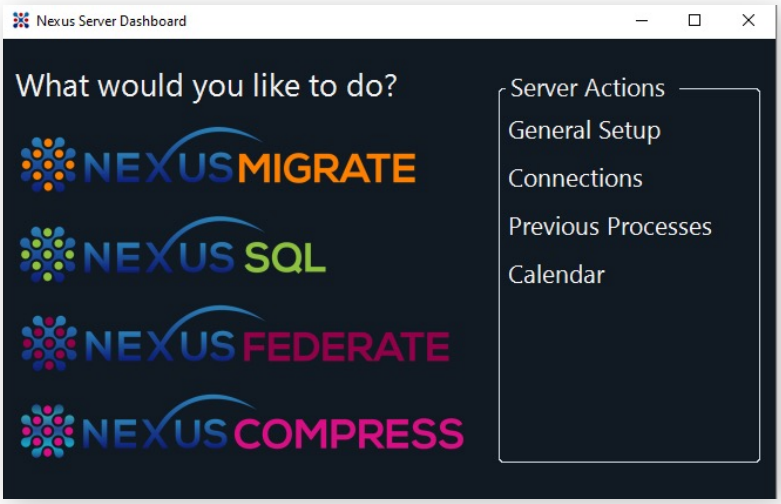

Simply choose your source and target systems and the tables to move.

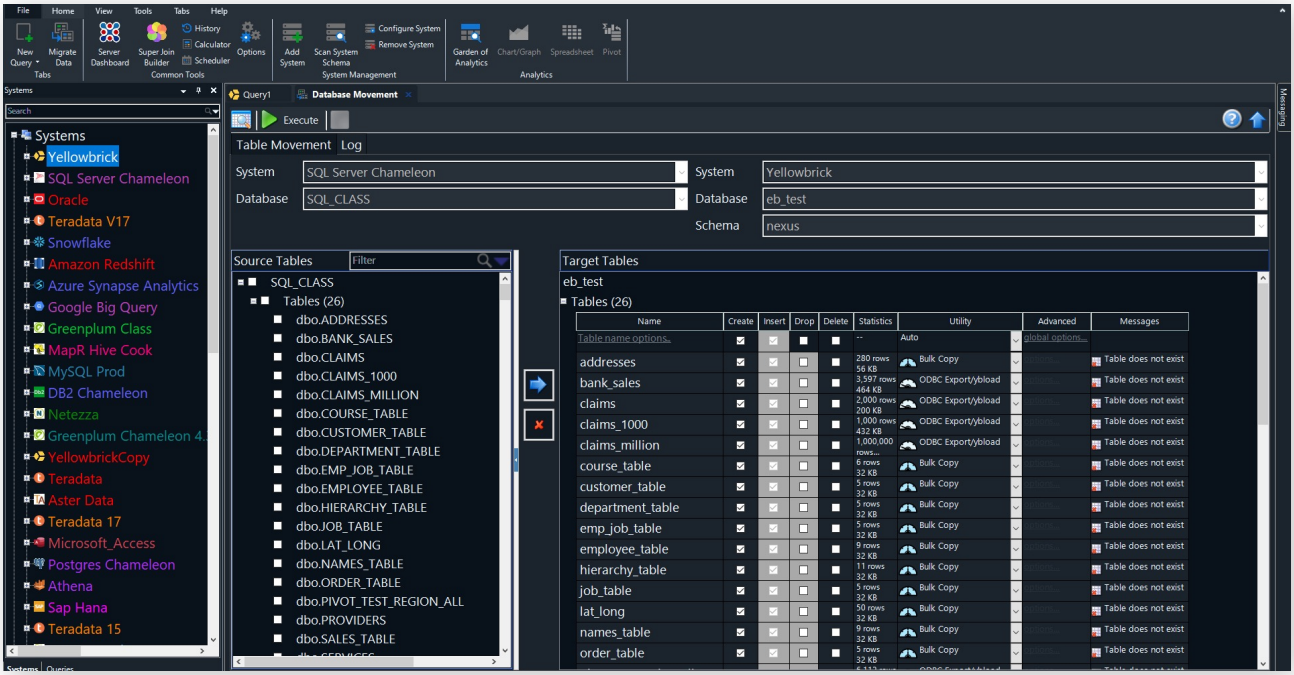

Nexus enables you to run queries on SQL Server and Yellowbrick side-by-side. Simultaneous queries provide key support in a POC.

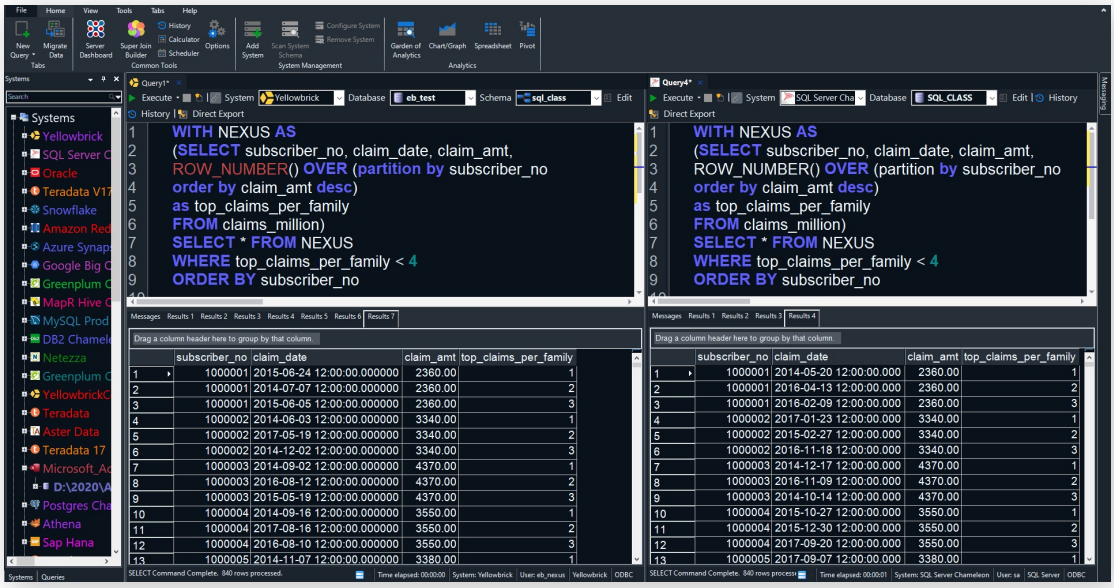

To run federated queries on SQL Server and Yellowbrick, simply place a checkmark on the columns you want on the report. Nexus builds the SQL, table conversions, and even allows you to choose where to process the join.

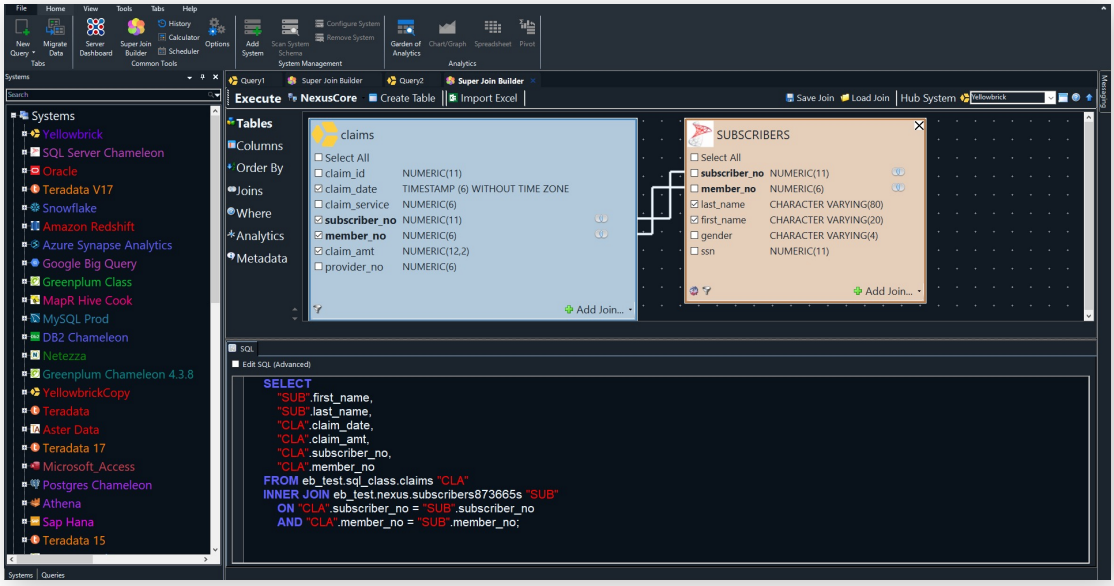

Summary

This guide explains why SQL Server is not suited for organizations looking to deliver a modern analytics experience and how Yellowbrick can provide a solution that runs 100X faster at a fraction of the cost. By partnering with Coffing Data Warehouse, Yellowbrick enables organizations to fast track the migration from SQL Server while reducing migration risk with the co-existence of both platforms.

To learn more about Yellowbrick Data, visit **yellowbrick.com** to book a demo or sign up for a free test drive.

About Coffing Data Warehousing

CoffingDW has been on the cutting edge of data warehousing technology for over the past 20 years. Our solutions have helped provide the largest data warehouses in the world increased cost savings, productivity, and efficiency. To learn more visit **CoffingDW.com**

COFFING
DATA WAREHOUSING