DATA ANALYTICS

The Associative Difference

Analytical power driven by
unique engine technology

LEAD WITH DATA™ Qlik

When a business user looks at a visualization or dashboard, they usually end up with more questions than answers. In order to make an informed decision, people have to ask and answer follow-up questions – going through a process of questioning, evaluating, and then questioning again – until the right insight is fully uncovered and understood. It's a process unique to every person and their line of thinking. And it shouldn't be limited to just data experts and analysts that can build visualizations and queries.

Almost all of the analytics tools on the market today are built on some form of query-based architecture. That's a major flaw. SQL and relational databases were was never designed to support the interactive analytical process, and they impose significant limits on exploration and discovery – including how data can be combined, how freely people can explore, and how fast calculations are made.


Tunnel Vision

x Partial subsets of data

x Restricted exploration

x Slow performance

With query-based BI tools, data sources are brought together using SQL joins, which can result in data loss and inaccuracy. Assumptions must be made in advance about what types of questions people will have, and there is little ability to pivot analysis to an adjacent question or data set. Users are limited to simple filtering on top of partial views of data – the query result set. New questions usually require new analytics to be built by experts. And as the system scales, performance will slow – especially when trying to calculate on large volumes of complex, record-level data.

The Qlik Associative Engine – Built for Interactive Analytics

Qlik offers an entirely different analytics technology: The Qlik Associative Engine. Designed specifically for interactive, free-form exploration and analysis, this powerful engine lets you make selections literally anywhere, and see the impacts of those selections in all analytics. You can see the relationships in your data, including what's related and unrelated to your selections. You can search across all your data to refine context. You can access all your data at all times, across all your data sources, from the highest-level summarization down to the full transaction-level detail. And with industry leading performance, calculations update as fast as you can think, no matter what selections you make, how complex the data is, or how many users there are. We call this The Associative Difference® – and only Qlik can deliver it.
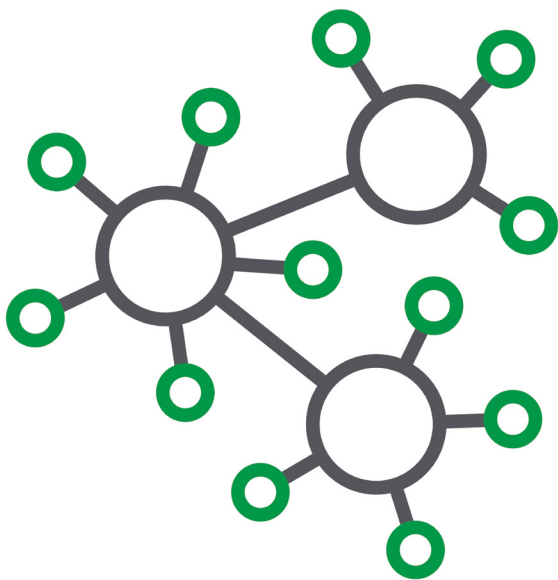
What makes Qlik's engine technology so powerful? Three major capabilities:

Data Combination
The ability to bring together all your data, from all your different sources, without leaving any data behind or getting anything wrong.

Interactive Exploration
Users of all skill levels can explore without boundaries, asking all the questions and follow-up questions they have, through interactive selections and search.

Speed
People can get answers as fast as they can think, prompting new questions and further exploration until insight is discovered.

Peripheral Vision

✓ All Your Data

✓ Explore Without Boundaries

✓ Speed of Thought
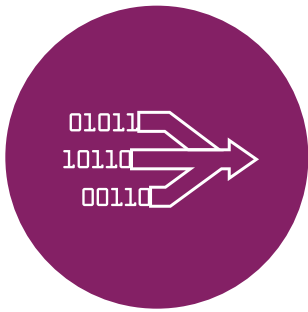
All Your Data

Query-based tools struggle when combining large numbers of data sources. Because of the way SQL joins work, they require primary and secondary data sources to be defined. This results in an intersection of data, and data can be left behind or mis-counted.

Qlik is different. Our associative engine fully combines all your data – even imperfect data – from large numbers of disparate sources, without suffering data loss or inaccuracies. It achieves the equivalent of what's technically known as a "many-to-many full outer join." Simply put, this full union allows users access all their data , from all their sources, at all times.

Query-Based Tools
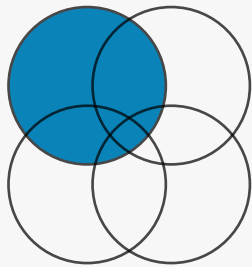
- You don't know what's been left behind.
- You don't know what has been counted or how often.

Qlik's Associative Engine
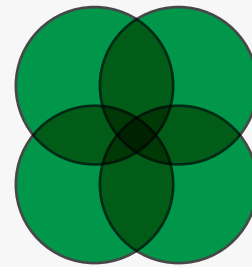
- No data is left behind.
- No data is mis-counted.

The Qlik Associative Engine allows imperfect data to be loaded without having to be fully cleansed and modeled in advance. This helps expose areas of your data and your business where issues may exist and value can be created. And users can create hierarchies on the fly, without having to pre-aggregate data in advance, in response to new questions.

Explore Without Boundaries

With Qlik's Associative Engine, people can explore and analyze freely, in any direction, without restrictions or boundaries. **Here's how:**

- Users can make selections within all visualizations, tables, charts, and other objects.

- Or they can use global search to refine context and select combinations of data.

- After each click, the engine instantly responds by
    - Recalculating all analytics to the new context (selection state)
    - Highlighting associations in the data values using green (selected), white (related), and gray (unrelated)

Because the engine dynamically calculates analytics instead of pre-aggregating data, users can pivot their thinking to new ideas or data sets. This means any user can ask any question, at any level of detail, without being limited by predefined queries or hierarchies.

Users also have the ability to see both related and unrelated values in their analysis, relative to their selections. Unrelated data often conveys the most impactful insights – products that didn't sell, for example, or customers who didn't buy – helping users hidden areas of opportunity or risk. With query-based tools, these values are simply filtered out, leaving people with only a partial data set and an incomplete perspective.
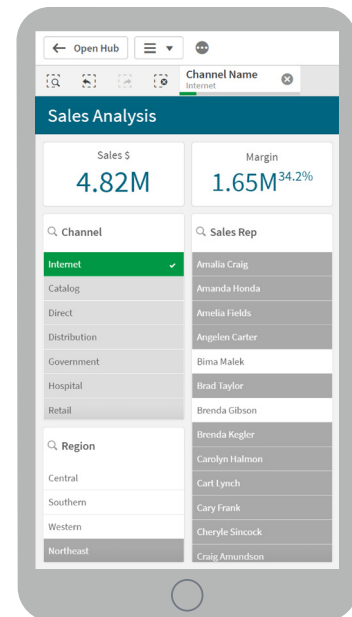
With Qlik, whenever a user makes a selection or performs a search, all visualizations, analytics, and associations are instantly updated to the new context. That's because the Qlik Associative Engine is context-aware, maintaining the selection state for all analytics across an entire application.

This "peripheral vision" gives people the ability to understand the impact of their questions on surrounding analytics, at different levels of detail, at the speed of thought. Users can immediately spot potential areas of interest, think of new questions, and continue to explore further. With query-based tools, you have to wire objects together and run multiple queries to simulate this experience, but it quickly breaks down. There is no central, context-aware engine.
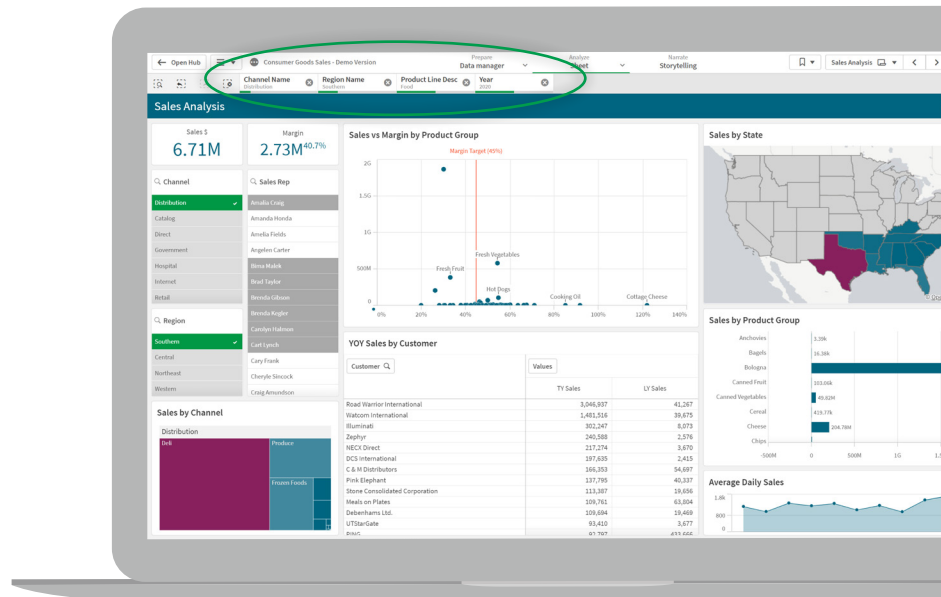
A leading global bank and investment firm discovered nearly $20 million in mortgage pipeline that wasn't associated with any particular loan processor. These mortgages showed up "in the gray." The firm immediately pursued the lost pipeline, generating tremendous value from a single discovery that wouldn't have been possible with query-based tools.

Speed of Thought

The Qlik Associative Engine is built for speed. It dynamically calculates analytics and highlights associations as quickly as users can think of questions. With instant responses to selections, people can get the answers they need without having to wait, allowing them to continue their thought process and explore further. And this means hundreds of more informed decisions every day.

In addition, query-based tools may be sufficient for creating visualizations, but only a handful of skilled users can or want to do this. So what happens to the majority of your business users? They end up with essentially static reports with limited filtering. And this leads to follow-up questions that can't be answered in real-time. So they have to go back to the data experts and wait. This "ask, wait, answer" cycle puts the brakes on getting insights.

Under the Hood

The Qlik Associative Engine achieves an unmatched combination of speed and flexibility that simply isn't possible with query-based tools. How? By combining several unique technology innovations – including compressed binary indexing, logical inference, and high-performance dynamic calculation.

Compressed Binary Indexing

Compressed Binary Indexing is the process that our engine uses to load and store data, creating a compressed, multi-table in-memory data store, with an associative index, optimized to support analytics calculation. This reduced, optimized data footprint is created by using binary pointers, storing each unique value only once and symbol pointers everywhere else. In addition, the engine maps all the relationships between the values in the data, resulting in a highly optimized data store supporting unmatched performance and scalability for analytics calculation.

This approach also provides a number of advantages for developers and organizations. First, due to binary compression, data loads extremely fast and has a smaller footprint in-memory. Second, due to the engine's ability to auto-join tables, content creators don't have to worry about joining tables in SQL statements when building analytics – it's done automatically, and therefore correctly.
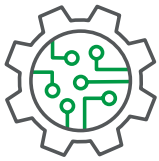
Logical Inference

Logical inference is the process in which the engine determines related and unrelated data values, relative to user selections, after each click. It leverages the mapping in the compressed binary index to determine what's related and what's not. It then flows the included data set into all analytics for dynamic calculation, and exposes related and unrelated values to the user in the green/white/gray color scheme.

This incredibly important process happens in the background without any additional effort required for developers. The benefits are all about context – as analytics are updated together automatically with no need to try and keep multiple queries in sync. And as described, the user gets unmatched exploratory power and insight, with the ability to make selections anywhere and see the impacts everywhere, in surrounding analytics and data relationships.

High-Performance Dynamic Calculation

High-performance Dynamic Calculation is the process of analytics calculation in our engine, based on a full set of record-level data, determined by logical inference. Analytics calculation is based on a full set of record-level data determined by logical inference, stored in a hypercube for each analytic. Our engine calculates at the lowest level of granularity with no pre-aggregation for maximum flexibility. It leverages advanced caching techniques to maximize speed for the overall process. Along with compressed binary storage optimized for calculation, and reduced (included) data sets determined by logical inference, calculation speed in Qlik Sense consistently outperforms the competition.

The result is a "speed of thought" workflow for users, where they can interact with data without having to wait, maximizing the potential of their thought process. And our engine maintains this level of performance when applications are scaled to large, concurrent users bases built on large complex data sets.

This analytics engine technology is our core advantage – and the muscle behind The Associative Difference. And just like a high-performance car, you don't need to be an engineer to feel the difference when you drive it.

Global tech giant Cisco brought together 500 million records of complex customer data from multiple systems and allowed their large sales force to freely explore customer portfolios. Users uncovered the best cross- sell opportunities for maintenance and additional products, helping the company generate $100 million in support renewals and $4 million in cost savings – and greatly improving customer satisfaction.
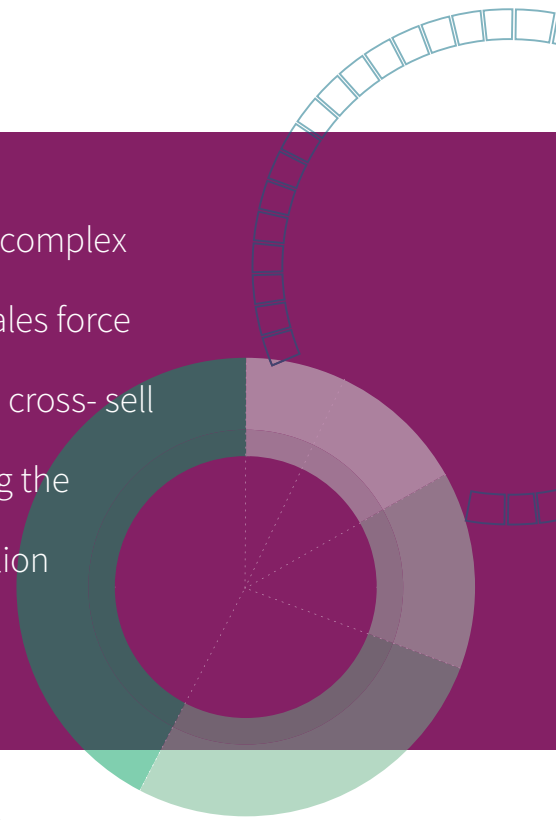
Qlik Q

About Qlik

Qlik's vision is a data-literate world, one where everyone can use data to improve decision-making and solve their most challenging problems. Only Qlik offers end-to-end, real-time data integration and analytics solutions that help organizations access and transform all their data into value. Qlik helps companies lead with data to see more deeply into customer behavior, reinvent business processes, discover new revenue streams, and balance risk and reward. Qlik does business in more than 100 countries and serves over 50,000 customers around the world.

qlik.com

ASSDIFF051021-CM