QlikView
ARCHITECTURAL OVERVIEW

A QlikView Technology White Paper
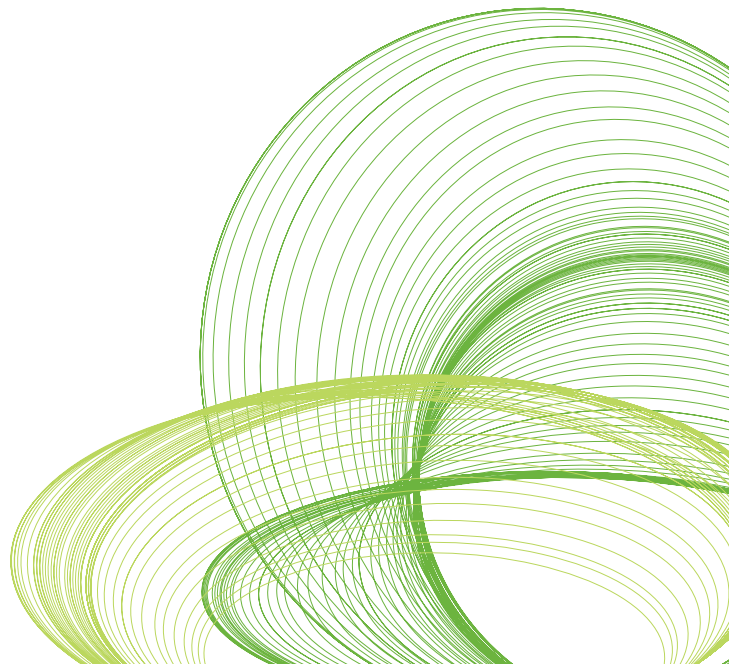
QlikView

Table of Contents

Making Sense of the QlikView Platform

Our customers often ask about what goes on under the hood of QlikView. In this *QlikView Technology White Paper* we show details of QlikView from two perspectives: components and functionality. Here you will find information about the relationship among QlikView components as well as the processes that enable QlikView to transform raw data into actionable analytics. The target audience for this paper consists of IT professionals and tech-savvy business people who are interested in business intelligence (BI) software.

This paper answers the question, "What are the basics I should know about the QlikView platform and its components so I can make an informed buying decision?" Its objectives are twofold: first, to help prospective and existing QlikView customers understand the components that comprise a QlikView deployment, and second, to help people unfamiliar with QlikView understand how data flows through a QlikView deployment on its way to becoming actionable information.
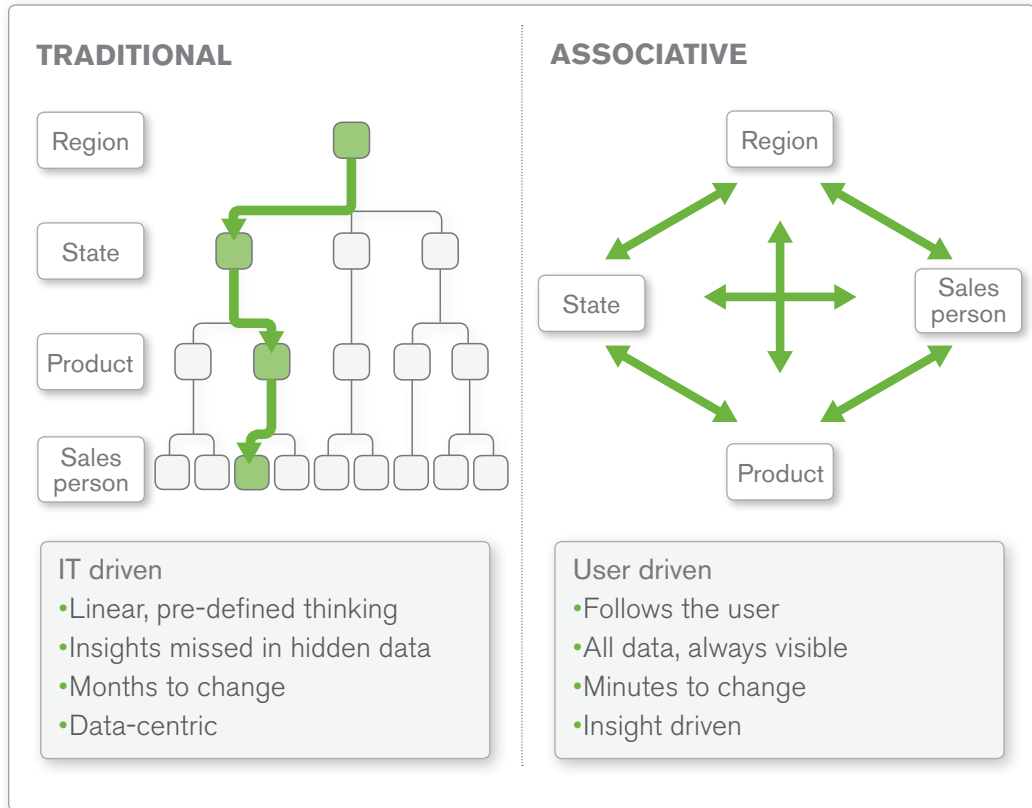
Most BI Software Is Built on Old Technology

Query-based BI tools have been the status quo for decision support for decades now. Many variations of query-based BI software are on the market. Some are flexible and others are high-performance. But they all share one critical flaw: they are unable to inherently maintain associations among data elements (see Figure 1).

- **Query-based tools divorce data from its context.** People making complex business decisions don't always have full access to their supporting data – even when they have access to BI software. Some data is available only as isolated and discrete queries, without context between one query and the next. This leaves gaps for people trying to make data-driven business decisions.

- **Where there is a query, there is dissociation.** With query-based tools, creating associations among all available data elements would require a business analyst or IT professional to cram every associated field into a single query — a nearly impossible task. The alternative — hard coding associations among queries into the application layer — is equally daunting.

Figure 1: Traditional BI software breaks data associations

TRADITIONAL

ASSOCIATIVE

Region
State
Product
Sales person

Region
State
Sales person
Product

IT driven
• Linear, pre-defined thinking
• Insights missed in hidden data
• Months to change
• Data-centric

User driven
• Follows the user
• All data, always visible
• Minutes to change
• Insight driven

## The QlikView Difference: The Associative Experience

QlikView is the world's first associative, in-memory business intelligence platform. QlikView manages associations among data sets at the engine level, not the application level, by storing individual tables in its in-memory associative engine. Every data point in the analytic dataset is associated with every other data point in the dataset. Datasets can be hundreds of tables with thousands of fields.

Unlike traditional query-based BI tools, when the QlikView user selects a data point, no queries are fired. Instead, all the other fields instantaneously filter and re-aggregate themselves based on the user's selection. Selections are highlighted in green. Datasets related to the selection are highlighted in white. Unrelated data is highlighted in gray. This provides a very intuitive, user-friendly way for people to navigate their data on their way to business insight. (For more detailed information about QlikView's associative experience, please see the *QlikView Technology White Paper*, "The Associative Experience: QlikView's Overwhelming Advantage.")

The QlikView Platform: A Component Overview

This component overview covers QlikView's individual software components, which support QlikView content creation (QlikView Developer), deployment (QlikView Server/Publisher), and consumption (QlikView clients) (see Figure 2).

Figure 2: QlikView — A component overview
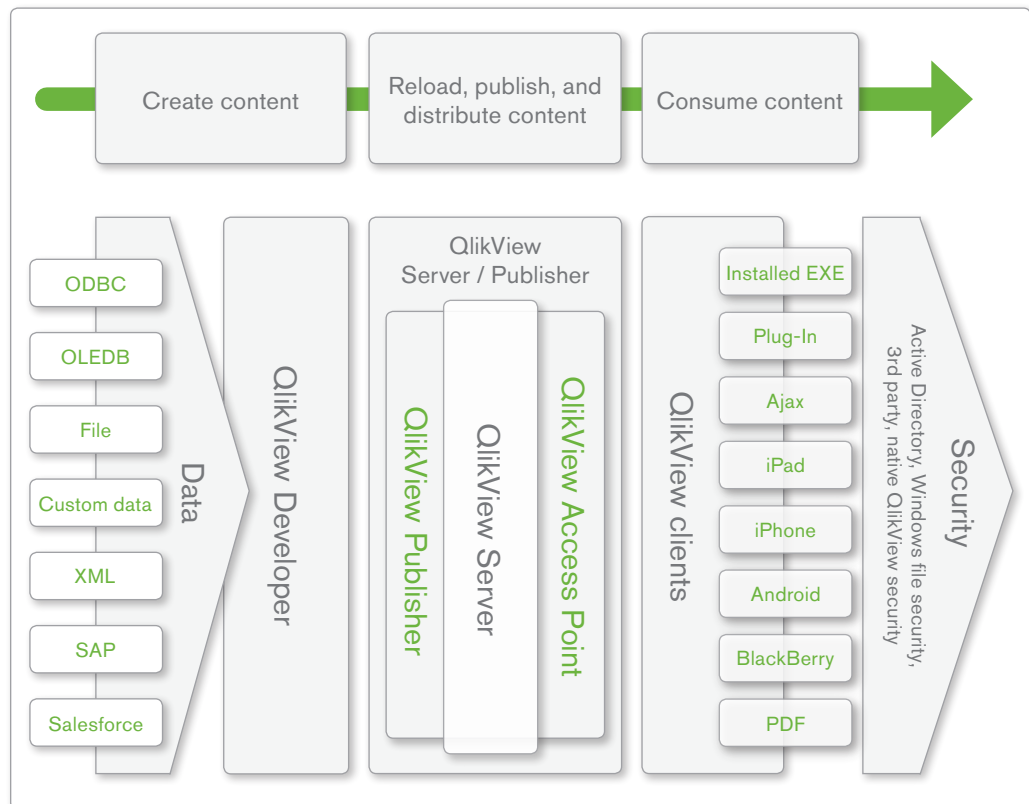
THE FIRST STEP: LOAD RAW DATA INTO QLIKVIEW

QlikView can extract and combine data from multiple, disparate sources without explicitly importing data from, or exporting it to, external systems. This enables centralized analysis of enterprise data regardless of where it originated, providing users with a holistic view of the business. QlikView can consolidate data from many different sources including:

- **ODBC- or OLEDB-compliant databases.** QlikView can use as a data source any database that is accessible via the open database connectivity (ODBC) or object linking and embedding database (OLEDB) standards. These include Oracle Database, Microsoft SQL Server, IBM DB2, MySQL, and others.

- **Standard file formats and systems that require specialized connectors.** In addition to common databases, QlikView can load data from a variety of standard file formats (e.g., HTML, Microsoft Excel, text, XML, etc.). QlikView also offers specialized connectors for reading data from SAP and Salesforce.com, as well as a custom data interface for extracting data from non-ODBC compliant databases.

GOT DATA? NOW IT'S TIME TO CREATE QLIKVIEW CONTENT

The QlikView Developer component enables developers to define the load script and visualizations for QlikView applications.

- **Load scripts define data sources.** QlikView load scripts define data sources as well as which data will be extracted from sources. These scripts also define what transformations, if any, should be applied to the data as it is loaded. While QlikView in and of itself is not a standalone extract, transform, and load (ETL) tool, QlikView's out-of-the-box ETL functionality is rich and powerful, with more than 350 functions for transforming data.

- **Visualizations turn numbers into pictures.** QlikView provides flexible, intuitive, and powerful data visualizations (see Figure 3). Once data has been extracted from source systems and transformed as needed, developers and designers use QlikView Developer to define which visualizations end users will ultimately interact with in the process of answering important business questions.

Figure 3: QlikView has a rich, visual user interface
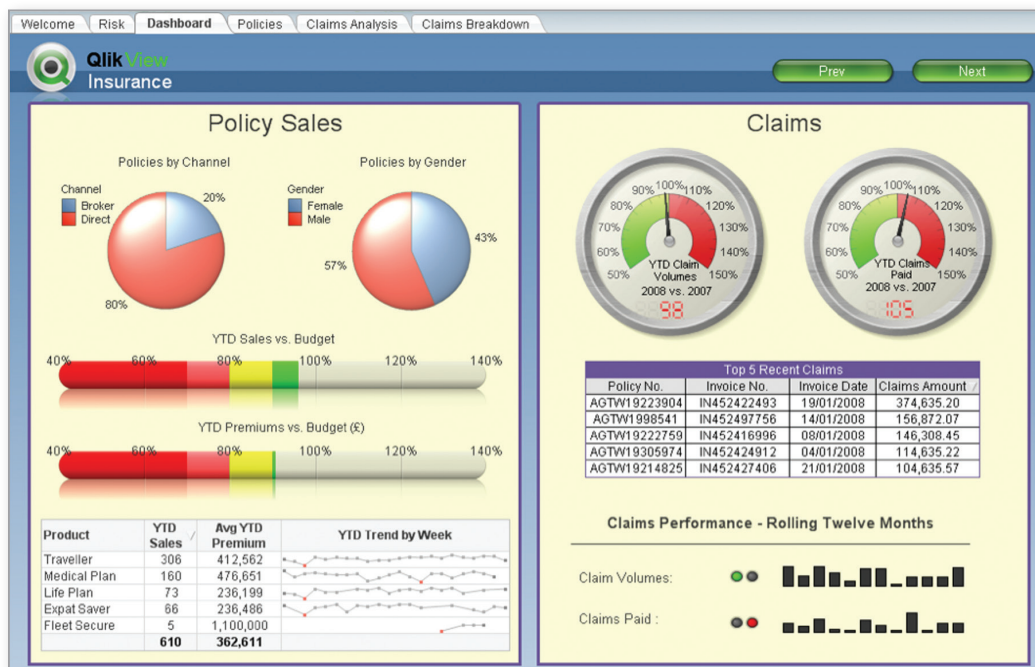
RELOAD, PUBLISH, AND DISTRIBUTE CONTENT

QlikView Server and its sub-components perform the tasks of reloading, securing, administering, and distributing QlikView content to end users.

- **All QlikView clients connect to a QlikView Server, where content is stored.** QlikView Server is responsible for providing client-server access to QlikView content. It stores QlikView documents and makes their content available to potentially thousands of end users. QlikView Server also helps with the scheduling of data reloads, though scheduling is typically handled by the powerful QlikView Publisher sub-component.

- **QlikView Publisher provides additional scheduling and distribution.** The optional QlikView Publisher component is designed to handle complex content deployment scenarios common in large enterprises. It extends and enhances QlikView Server's scheduling functionality and provides additional security for QlikView content based on users and groups. For example, from a central QlikView document, Publisher can create separate, individualized sub-documents for a user or group. These documents contain only the needed or desired data.

- **AccessPoint is a customizable entry point to all QlikView applications.** Via the AccessPoint sub-component of QlikView Server, end users can access all QlikView content they are authorized to view. In addition, QlikView AccessPoint provides back-end services such as load balancing user sessions among multiple QlikView Servers in a clustered QlikView Server deployment.

USERS CONSUME CONTENT ANY TIME, ANYWHERE

Once QlikView content has been deployed via QlikView Server, it is ready to be consumed by end users. One of QlikTech's central tenets is that QlikView content should be available to the end user where and when it is needed. To that end, QlikView supports a variety of technologies and ways to connect to the QlikView Server.

- **Browser access lets users get to their data from any PC or laptop.** QlikView offers two browser based clients: a Microsoft Internet Explorer plug-in and an Ajax client. The Ajax client is a zero-install option.

- **Mobile device support means data goes wherever users go.** The popularity of portable devices such as smart phones and tablets will only increase over time. QlikView supports leading mobile platforms such as Apple iPhone and iPad, BlackBerry, and Android devices.

- **The installed client meets the need for offline access.** For cases where online access is not possible, the self-contained nature of QlikView documents means that users can analyze data locally in an offline mode using the installed client. The installed client can also be used in a connected, client-server mode if desired.

- **Emailed reports let people live in their email if they want to.** Business users who are satisfied with static content can consume QlikView reports as PDF documents. The QlikView Publisher sub-component of QlikView Server handles distribution of PDFs to users via email.

A Look at The QlikView Platform From a Functional Perspective

A functional perspective on QlikView's core processes can help IT professionals and tech-savvy decision makers better understand what makes QlikView so special (see Figure 4).

Figure 4: QlikView — A functional perspective
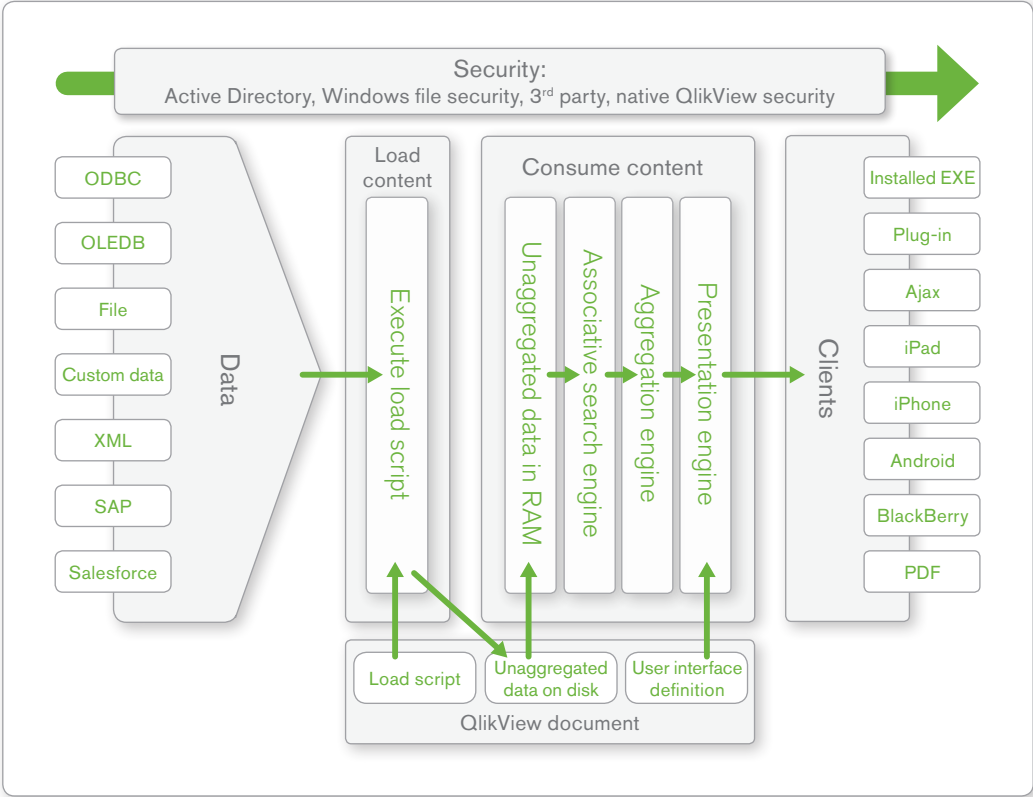
THE QLIKVIEW FILE CONTAINS EVERYTHING NEEDED TO SUPPORT QLIKVIEW ANALYTICS

The native QlikView file format, also known as the QlikView document, is self-contained, portable, and designed for ease of use. QlikView documents contain:

• **The load script.** QlikView load scripts extract data and process the unaggregated data that has been returned from the load process. QlikView documents store data in a highly compressed format.

• **User interface definitions for visualizations.** QlikView documents contain definitions for the charts, graphs, and other dashboard and reporting elements of QlikView.

• **The actual data to be analyzed.** QlikView documents can also contain the data to be analyzed in a highly compressed format. This enables both offline and client/server-based analysis by loading the data into memory as needed.

QLIKVIEW DOCUMENTS SERVE UP DATA STORED ON A QLIKVIEW SERVER

When a QlikView document is published on a QlikView Server, the content it contains becomes available for consumption by any end user with privileges to access it.

• **When a user first opens a QlikView document, data is loaded in memory.** The compressed and unaggregated dataset is loaded off the disk and into the QlikView Server's random access memory (RAM). This in-memory repository serves as the base dataset for this initial user and all other users requesting the same document. This repository stays in memory until no user activity has occurred within a defined time-out period.

• **Users explore data via selections.** Central to QlikView is the concept of a user-defined selection state. As users click around in a QlikView document, they indicate which subsets of data they are interested in analyzing and which subsets should be ignored. QlikView takes advantage of the highly indexed nature of the unaggregated dataset. It dynamically presents a subset of all the data available to the QlikView document based on the selection state. This happens in real time as the user executes clicks.

• **Upon selection, aggregates render instantly.** QlikView renders aggregates as intuitive and interactive user interface objects: charts, graphs, tables, etc. Users interact with objects in QlikView documents through any supported client. End users can create their own objects using the collaboration features of QlikView.

QlikView's Architecture Raises the Bar

Through the technology decisions we've made, QlikTech gives our customers the fastest, most intuitive, most flexible BI platform in history. The QlikView platform delivers:

- **A rich, full, overview of the business.** With QlikView, all user interface objects and their underlying aggregates are always in context (associated) with one another. This associative experience gives decision makers a better overview of their business. They can create and use dimensions in charts, and consume content any way they want to: an installed client that works offline, a Web browser, or a mobile device. Business data is available any place, anytime.

- **Broad, deep business insights.** Nearly all organizations deal with large datasets. Loading large volumes of data into QlikView from a combination of many data sources allows insights into aspects of the organization not otherwise apparent. With QlikView, even large volumes of data don't have to be pre-aggregated. Users get near-instantaneous response speeds as each and every QlikView object responds to their every click.

- **A shared version of the truth for all.** QlikView can scale to support thousands of connected users, giving everyone in the organization access to one shared version of the truth. With QlikView, an application originally intended for a single user or small group can be deployed to many others simply by adding memory and processing power. The application itself does not have to be changed.

- **Speedy time to value.** QlikView is quickly and easily deployed and integrated with existing enterprise systems. Developers don't have to spend time hard-coding answers to specific business questions. This saves a vast amount of time in the implementation process. As a result, users get answers to questions when they need them, rather than waiting weeks or months. High user adoption and satisfaction invariably means better return on investment.

- **Low ongoing total cost of ownership.** QlikView offers a simple way to create and manage data analysis. QlikView simplifies the extraction and combining of various data sources to create a single, detailed, flexible view of an organization's data. The platform has a central, web-based management console that enables administrators, wherever they may be, to manage everything that has to do with QlikView.

Appendix

Links to Related QlikView Technology White Papers

The Associative Experience: QlikView's Overwhelming Advantage, October 2010

Links to Related QlikView Blog Posts

"Unpredictable Questions and the Power of Gray," August 18, 2010

http://community.qlikview.com/blogs/theqlikviewblog/archive/2010/08/18/
unpredictable-questions-and-the-power-of-gray.aspx

"QlikView Is Associative to Its Very Core," August 16, 2010

http://community.qlikview.com/blogs/theqlikviewblog/archive/2010/08/16/
qlikview-is-associative-at-its-very-core.aspx

"The Car Engine Analogy," August 13, 2010

http://community.qlikview.com/blogs/theqlikviewblog/archive/2010/08/13/
car-engines-and-the-associative-experience-an-analogy.aspx

"It All Comes Down to Simplicity," August 12, 2010

http://community.qlikview.com/blogs/theqlikviewblog/archive/2010/08/12/
simplicity-equals-success.aspx

"QlikView Users Have an Emotional Attachment to the Associative Experience,"
August 6, 2010

http://community.qlikview.com/blogs/theqlikviewblog/archive/2010/08/06/
qlikview-users-have-an-emotional-attachment-to-the-associative-experience.aspx

"The Consumer Enterprise: I Want My Home Technology at Work," August 3, 2010

http://community.qlikview.com/blogs/theqlikviewblog/archive/2010/08/03/
the-consumer-enterprise-i-want-my-home-technology-experience-at-work.aspx