

Qlik[®] Sense architectural overview

March, 2015



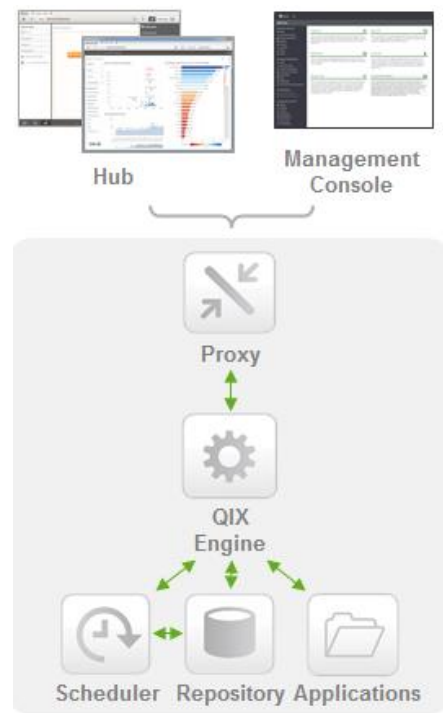
Platform

Qlik® Sense is an analytics platform powered by an associative, in-memory analytics engine. Based on users' selections, calculations are computed at runtime against data stored in-memory. Results are returned to users via a zero footprint web interface delivered on desktops, laptops, mobile devices, and through embedded analytics. Qlik Sense offers a highly interactive, associative experience in which users can freely navigate through data with little to no constraint in their analysis path.

Overview

Qlik Sense is designed to support self-service visualization in a scalable, secure, and governable way. Qlik Sense can be deployed on a single server and scale vertically and horizontally to address the availability and processing requirements of your deployment, whether on premise or in the cloud. The server-side components are described below. The client interface, called the Hub, and the administrative interface, called the Management Console, are described in the section titled Qlik Sense Clients below.

- **Qlik Sense Proxy (QPS)** is the entry point for users and administrators via the Hub and Management Console, respectively. This component integrates with many identity providers (e.g., Active Directory). The Proxy also manages sessions, provisions licenses, and handles load balancing to the other components.
- **Qlik Sense Engine (QIX)** is the associative, in-memory data indexing engine. This RAM-based application tier delivers highly interactive self-service visualizations, search, and calculations at runtime.
- **Qlik Sense Scheduler (QSS)** is the scheduling component that coordinates data loads. This scheduler supports time based application reloads, dependencies, and may be driven by enterprise schedulers (e.g., CA AutoSys).
- **Qlik Sense Repository (QRS)** is the centralized store of configuration and management information of the platform. This component manages user definitions, security, and many other elements of the platform.
- **Qlik Sense Applications (.QVF)** are comprised of data, a data model, and presentation layer. These applications are persistently stored on a file system and are loaded into memory by the Qlik Sense Engine as users request them.



Qlik Sense Architecture

Associative, in-memory applications

Qlik's QIX Engine uses an associative, in-memory technology that enables users to analyze and freely navigate data in an intuitive way. This second generation of Qlik's proven data indexing engine allows users to freely explore data and create visualizations based on data from multiple data sources simultaneously, ranging from Excel and Access, to databases such as Oracle and SQL Server, to big data sources such as Hadoop and Redshift.

Qlik Sense uses columnar, in-memory storage. Unique entries are only stored once in-memory and relationships among data elements are represented as pointers. This allows for significant data compression, more data in RAM, and faster response times for users.

There are some big data scenarios where data should remain at the source. Using Direct Discovery, data sources can be queried directly as the user clicks, yet still provide an associative experience to the user.

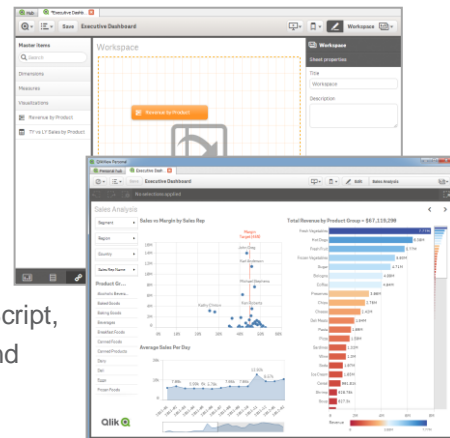
What does "associative" mean?

Associative refers to the unique combination of in-memory data storage and a patented engine. This allows users to search and navigate through and create visualizations upon data with a **continuous context**. This removes the constraints imposed by traditional hierarchical or query based approaches.

Qlik Sense Clients

Hub

The Hub is the Qlik Sense user environment is delivered through a zero-footprint web browser interface. Through the Hub, all aspects of development, drag-and-drop content creation, and consumption is possible. Qlik Sense delivers a *responsive design* methodology to automatically display and resize visualizations with the appropriate layout and information to fit the device, whether it a browser on a laptop or desktop, tablet, or smartphone. Built with current standards of HTML5, CSS3, JavaScript, and web sockets, Qlik Sense enables you build on any platform and deploy to any platform.



The Hub

Management Console

The Qlik Sense Management Console allows administrators to govern all aspects of the Qlik Sense platform, ranging from data connectivity, application and task management, to security administration, monitoring, and auditing.

Deployment Console

The Qlik Deployment Console allows administrators to configure and monitor Qlik Sense environments. Predefined templates for VMWare vSphere, Amazon AWS, and physical hosts allow for rapid creation, scaling, and monitoring of Qlik Sense environments. Migration tools within this console allow for content migration among environments to support traditional development-test-production development lifecycles.

Distributed Architecture

Qlik® Sense has a modular, distributed architecture designed to meet the unique requirements of your organization. While a single Qlik Sense server support hundreds of concurrent users, Qlik Sense can scale even further with its distributed architecture of independent nodes.

Distributed Workloads

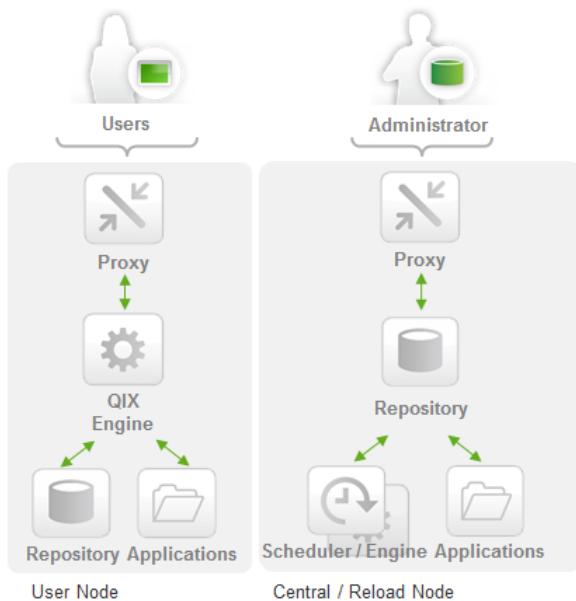
The Qlik Sense components can be deployed on servers in different combinations to accomplish specific functions. The table below describes how the components of Qlik Sense are deployed to allow servers to function as User nodes, Reload nodes, and the Central node.

Function	Description	Component				Comment
		Proxy	Engine	Scheduler	Repository	
User Node	Self-Service Analytics for Users	X	X		X	Proxy can be separated to its own server.
Reload Node	Application and Data Reloads		X	X	X	Engine performs reloads invoked by scheduler.
Central Node	Primary Repository and Application Store	X			X	Proxy used for administrator access.

The repository component is mandatory on each server and each is constantly synchronized to maintain consistent configuration and license information throughout the platform. Every Qlik Sense environment has a Central Node that acts as the primary repository and application store for the platform.

Basic Two-Server Deployment

This two node deployment example separates the workload generated by users performing analysis from data loads managed by the scheduler.



Synchronization

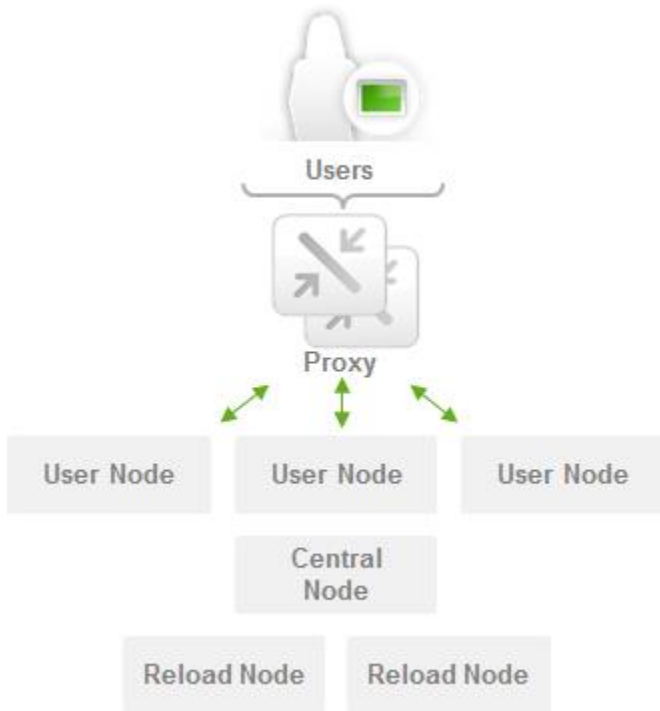
Repository databases are kept in sync via a transaction replication process. This lightweight sync process is multi-directional ensuring that information added to any node makes it to all the others.

Applications are synchronized only to those nodes that need them. UI changes are synchronized separately from data.

Synchronization processes provide for resilience within the Qlik Sense platform.

Scaling horizontally – Multiple servers

Qlik® Sense is able to scale workloads horizontally. Multiple Proxies, User Nodes for self-service analytics, and Reload Nodes for data loads enable Qlik Sense to scale horizontally for resilience and a growing workload.



About Workloads

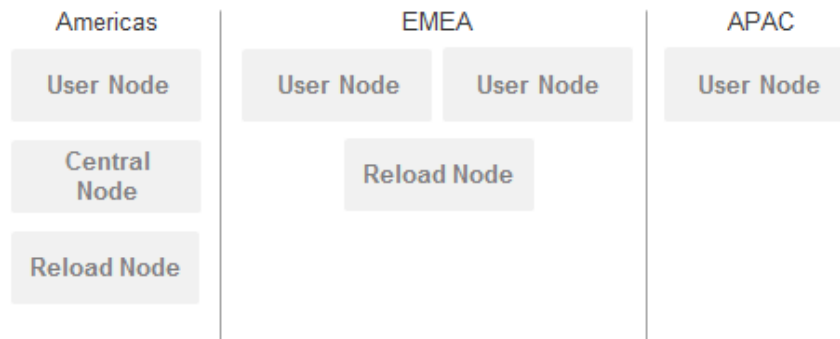
Security Rules are configurable and further enable you to dedicate workloads to specific nodes. For example, within **User Nodes**, isolate content creation from consumption, or segment applications by business unit. Within **Reload Nodes**, isolate data layer creation from application creation, or segment data loads by complexity or duration.

This completely configurable capability is discussed further in the document titled “Qlik Security Overview Whitepaper”.

Global Deployments



Qlik Sense is deployable at a truly global scale. The distributed architecture and synchronization allows a single Qlik Sense environment to be deployed across regions and domains with centralized administration and shared licensing. It is possible to deploy different server configurations to each region to support users’ specific self-service requirements.



Qlik® Sense APIs

Many organizations are tasked with integrating their platforms into an ever more complex and changing enterprise architecture. The Qlik Sense APIs ensure that Qlik Sense can integrate seamlessly into your environment to provide the security, automation, and delivery capabilities that are required.

The Qlik Sense APIs are REST and JSON based and are the *same* APIs used to develop Qlik Sense itself. This gives you tremendous power to build your own applications on the Qlik Sense platform. You can extend Qlik Sense by adding your own visualizations and your own data sources. You can leverage your existing security providers. You can integrate Qlik Sense with web applications or with desktop applications based on the .NET framework. You can manage licensing, users, and tasks programmatically.

Qlik Sense Extensions API

- Create New Visualizations

Qlik Sense Mashups API / .NET SDK

- Embed Visualizations and Functionality

Qlik Sense QVX SDK

- Create Custom Data Connectors

Qlik Sense Repository API

- Management and Administration

Qlik Sense Proxy API

- Authentication, Load Balancing, Session Management

Qlik Sense Directory Service API

- Create Custom User Directory Connectors

Qlik Sense Engine / Backend API

- Communicate Directly with Associative Engine

See the data sheet titled “Qlik Sense APIs” for more information.

To learn more, visit qlik.com.