Qlik Sense Topologies

Topology design principles

Dashboard (App)

QVF — A Qlik Sense app is a collection of reusable data items (measures, dimensions and visualizations), sheets and stories. It is a self-contained entity that includes the script, visualization and data to analyze in a structured data model.

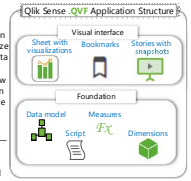

QVD — a Qlikview Data file(QVD) can store a single table of data for re-use purposes.

Data integration — the ability for cleansing, transforming, and unifying multiple, disparate data sources for analysis, without requiring external tools or data repositories. This includes databases, web content, and big data sources. (Similar to the activities which are normally done in a datawarehouse)

Direct Discovery – Provides a "hybrid mode" by loading the dimensions in the associative in-memory model. Combine this model with a real time SQL query connection to the 'Big Data' fact table. Providing an associative experience combined with real time facts.

QIX Accociative Data Index Engine Overview

The QIX engine is the second generation of the proven and patented engine technology that has powered Qlik products for the last decade.

Using Qlik Associative Data Indexing and dynamic calculation that empowers people to naturally explore using their intuition. The QIX engine exposes relationships in complex, multi-source data sets that would otherwise be hidden in hierarchical or query based approaches

Users benefit from QIX Associative engine every time they make a selection or search, allowing them to gain understanding and take next steps.

Important The QIX Associative Data Indexing Engine processing component in Qlik utilizes the CPU and memory from the sErver. The Qlik Sense architecture provides application orchestration, allows utilization of specific apps onto dedicated consume node. A combination Of Qlik Sense roles is possible but only if there are enough system resources available. Running both of these roles on a single virtual machine requires a 6-8-core VM.

A Qlik Sense Site consists of multiple servers. It is a distributed architecture that consists of one or more nodes (that is, server machines) that together form a site. By configuring which services to run on a node, it can be set up to perform a specific role within a site.

A server can contain all roles or configured for a specific Qlik Sense role. The following four roles are available.

Qlik Sense Roles

1. Proxy / Hub server

Contains the user interface called Qlik Sense Hub. It handles requests through SSL/TLS and WebSockets.

- Acts as a webserver
- Entry point into Qlik Sense for users and Administrators.
- Entry point for external system integration by using the mashup, authentication, session and load balancing API's

2. Consume Server Roles QIX / QES

Provides the Engine service features of Qlik core technologies. Use the Engine on Qlik Management Console to assign Engine services to specific QIX engine application servers.

- In-Memory, Associative Data Indexing Engine
- All engine services will utilize three Services
- Qlik Engine, Qlik Repository Services and Qlik Repository Database
- Identically configured application servers can be used for application orchestrion.
- Application nodes, look for services that consume a disproportionate amount of resources and consider placing these services on dedicated hardware.

3. Qlik Central Node

Controls the whole Qlik Site.

- Each role on the Central node is considered to be the master and contains information about the entire site. All nodes connect with this node.
- The QMC always uses the Qlik Sense Proxy Service (QPS) on the central node

4. Qlik Scheduler role (QSS)

Manage the scheduled tasks

- Handles application reloads
- Handles task chaining
- To which slave the QSS distributes the task ID to is determined by a load balancing operation performed by the master QSS on the master QSS

Qlik Repository Service & Database

Repository Services — Manage multi-node database communications. Each node requires a QRS and keeps a complete repository and transaction log. The repository contains information such as configuration, usage and rules. In a Qlik Sense environment, all Meta data is stored in a PostGreSQL Database. The database can be deployed to multiple servers.

In a multi-node environment, group databases by roles and deploy these to multiple database servers.

Legend for database icons

Entity (metadata)
The repository database contains the system configuration and all meta data about apps. The repository database is controlled by the Qlik Sense Repository Service (QRS) and stored in the Repository Database (QRD)
Default Path =%ProgramData%\Qlik\Sense\Repository\PostgreSQL\9.3

Windows Services

Qlik Sense's services are the cornerstones of the Qlik Sense architecture and is configured to support a variety of enterprise-ready deployment scenarios, including those that are geographically disbursed and with multiple data centers. A node in a Qlik Sense site runs a set of Qlik Sense services.
By configuring which windows services to run on a node, it can be set up to perform a specific role (for example, as a proxy node or a reload node) within a site. Deployment of these components requires planning.

Qlik Sense Node

Servers in a site are referred as "RIM nodes"or slave nodes. A node in a Qlik Sense site runs a set of Qlik Sense services. A multi-node environment incorporates a minimum of one server. One node assumes the role of central node, which is used as the (Master)central point of control.

Node synchronization is based on a peer to peer concept. The P2P system connects all nodes in a Qlik Sense Site. The configured nodes dynamically participate in traffic routing of apps in small segments to the designated nodes. When a change is made on a node, the resulting transactions are recorded in a transaction log. During the synchronization, the latest set of transactions from the log is sent to the other nodes and replayed. Applications synchronization to the RIM Nodes does not include QVDs. The central node requests updates from the other nodes with an interval of 15 seconds. Each of the other nodes initiates a synchronization session with the central node every 15 seconds. Between nodes asynchronous communication is being used.

Binary Sync (application)
The app data files contain the data models and app definitions. The app data files are controlled by the Qlik Sense Engine Service (QES).
Default Path =%ProgramData%\Qlik\Sense\Apps

Streams

The apps are organized in streams. A stream is a collection of apps that a specific group of users have access to.
The users of the stream can have different access rights.
Users might have read only access for the apps published to a stream, referred to as *Consumer*, or they may also have publish rights, allowing them to publish their apps to the stream, referred to as *Contributor*.

By default, Qlik Sense includes a stream called Everyone, in which all users have both read and publish rights to.

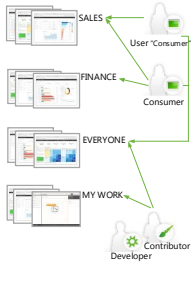

Central Governed Libraries

Use of centrally governed libraries for dimensions, measures and charts
A shared, centralized repository for data and visualization definition within a Qlik Sense app. Libraries provide a way for those responsible for managing the data to make it available to others in a way that makes it easier to use. But libraries include more than just data. Measures can be added to the library to ensure that everyone has a common definition for calculated results. And, even visualizations can be added to the library for easy reuse by others.

Without compromising the organisation's IT security regulations, users can build their own visualizations from a centralized library of pre-built data sets, expressions, and visualizations to ensure consistent use of data and values.

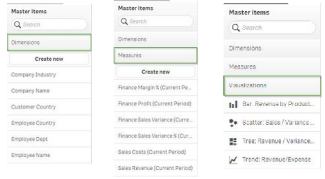

Select chart type, dimension and the measure. (Definitions "can not" be changed)

Governed Data Extraction, Transformation and Load Process (ETL)

ETL is a standard process which describes the moving and transformation of data from multiple data sources to more comprehensive views into a final sort of "star schema" model. By splitting the process into parts, roles and responsibilities can be divided and each layer can have it's own Sense Dashboard to monitor correctness of data and processes.

key features

- IT maintains and monitors the data model
- Master data stewards ensure correctness of master data
- Business key users can maintain business KPI's per department
- Dashboard design is splitted from data modelling
- End users have self service BI without changing KPI's or measures

Data Loading

Extraction QVF
- Without modifications, make copy of data from source systems (load * from tables)
- Store unedited separate tables in Qlik compressed QVD files

Transformation QVF
- Fast load of data from QVD's
- Align primary and foreign keys
- Cast data to correct data format (date/number/text/..)
- Create entities like customer, materials, sales facts
- Store into QVD for re-use.

Enrich master data QVF
- Ensure 1 version of truth of meta data
- Enrich data with business KPI's
- Add extra master data (segments/clusters/hierarchies) needed by business which are not in source

Create business model QVF
- Create model per business area
- Create star or snowflake schema data model to be used by presentation apps.

Dashboard QVF
- Users consume visualization apps to quickly answer the next question
- Users create new sheets using the pre-built *read only* Library with Dimensions, Measures and visualizations for "own use".
- Key users publish content into "community".
- Power users can load extra data and create new measures and charts

Data Governance

Verify data quality per table. *Customers without name*

Verify relationships between table/ entities. *Invoices without customer*

Define business rules for process analysis. (e.g. 3 way match: Invoice without goods receipt)

Governed Self Service Analytics

Enterprise Class Architecture

Qlik Sense Single Server (Node)

Qlik Sense's singular architecture leverages the latest web specifications, data interchange formats, and protocols such as HTML5, CSS3, JSON and WebSockets.

- WebSockets protocol provides full-duplex communication between the client and server over a single TCP connection.
- Users of all types can easily create and analyze in Qlik Sense using a unified HTML5 web based client, with no necessity for browser plugins, apps or fat desktop clients.
- IT can utilize a simple and powerful, web based management console (QMC) for administration developed with the same technologies as the front-end client.
- Developers can integrate data and build analytics using the web client, and can extend and customize Qlik Sense through standard and open API's

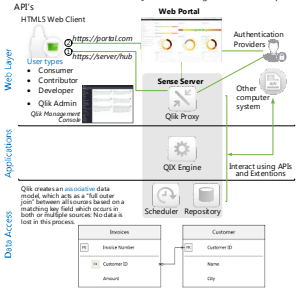

Multi-Layer Security

Qlik Sense has multiple layers of security (both out of the box and custom) that all work together to conform to the organization's security needs, resulting in A solid security model. Qlik Sense utilizes the following infrastructure components:

Network security
All communications that build trust between Qlik Sense services and clients are based on web protocols such as Secure Socket Layer (SSL) and Transport Layer Security (TLS). These protocols handle encryption and The exchange of information, keys and authentication certificates.

Server security
Qlik Sense uses the server's operating system security layer to control and protect Qlik Sense resources (files, memory, processes, and certificates) on the server.

Application security
Combined with the security that Qlik Sense provides – authentication, rules based content security, and dynamic data reduction, the result is an integrated, flexible and robust security model for Qlik Sense Security – This layer is the core of the platform's protection, comprised of authentication, authorization, auditing, confidentiality and availability.

Continuous security test and enhancement
This layer of the security model focuses on ensuring that the Qlik software is thoroughly analyzed from a security perspective by using *rugged* development practices as well as threat analysis and exploratory security and penetration testing. - Being "Rugged" – is about staying ahead of a threat. Using experimental approach to learn and improve while actively seeking out threats and creating defenses.

Qlik Sense Multi-Node

A common best practice for scalability, performance or security reasons is to create a three-tier model; web, application and data access. You can start with one node and increase nodes based on requirements to create a high available and redundant clustered environment.
Each node operates independently, which increases the system resilience, reduces maintenance, and increases the deployment flexibility.

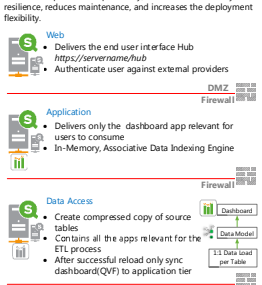

Legend for app icons

Node contains apps, serves dashboards to the users

Node contains apps but serves no dashboards to the users

Scalability

Virtual Topologies

Virtualized topologies are supported. Application behaviour depends on the capacity of physical hosts, desired ratio or virtual machines to hosts, and the underlying virtualization technology. This design model illustrates a fully fault-tolerant, virtual environment (VMs) and a plan for scaling out different Qlik Sense roles.

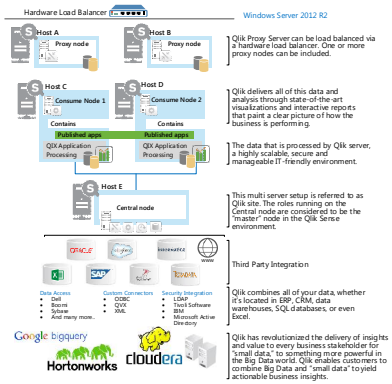

Qlik Sense Communication Flow

This scenario describes the traffic flow between multiple nodes.

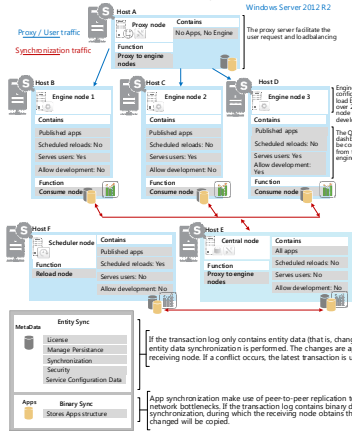

Installation concept

When you install multiple engines and proxies, the proxies load balance users randomly across all engines by default. This can be modified so that a proxy only talks to its local engine or to a sub-set of the engines, which caters for a number of deployment options to support various scenarios.

In this deployment example, the Qlik Sense site consists of the following nodes:

NODE NAME / SERVER ROLE	REPOSITORY SERVICE	QIX ENGINE SERVICE	SCHEDULER SERVICE	PROXY SERVICE
Central Node	✓	✓	✓	✓
Scheduler Node	✓	✓	✓	
Consume Node	✓	✓		
Proxy Node	✓			

In very large environments roles can be directed to dedicated server nodes.

The table above lists the Qlik Sense services that are deployed on each node in this site. Each node requires a QRS and QRD

A Qlik (application) environment separates the query processing component and indexing component to dedicated servers. The remaining roles can reside on two all-purpose application servers. on two all-purpose application servers.

Hybrid Solution - *Qlikview and Qlik Sense*

Hybrid Solution

This hybrid solution allows customers to reuse information generated by QlikView guided analytics. Qlik Sense can read Qlikview QVDs and QVWs (binary load) into Qlik Sense to be used for self-service visualization users.

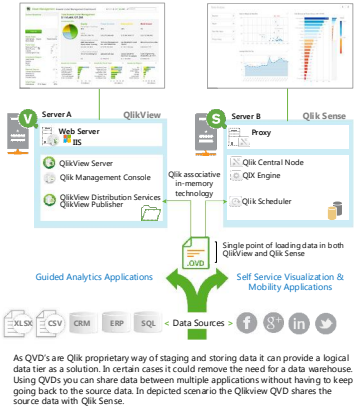

As QVD's are Qlik proprietary way of staging and storing data it can provide a logical data tier as a solution. In certain cases it could remove the need for a data warehouse. Using QVDs you can share data between multiple applications without having to keep going back to the source data. In depicted scenario the Qlikview QVD shares the source data with Qlik Sense.

Qlik Sense Governance

Risk and Controls

In order to provide a reliable and secure information management process the following key risks need to be mitigated:

- Unreliable reporting
- Multiple versions of the truth
- Users requirements are not met, causing additional decentralized and ungoverned tools (like Excel) to be used, e.g. not sufficient Self Service capabilities
- Performance issues
- Information is disclosed to unauthorized persons

An integrated (business and IT) governance framework allows you Know the risks that can cause unreliable reporting
Define controls to prevent or limit the impact of those risks
Define a persons with clear roles and responsibilities to ensure everyone only performs the activities for which they are trained/authorized.

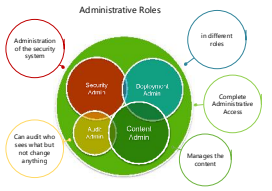

Governed Self Service BI

In the past, IT would create a reporting environment and the business users would just be able to read what IT delivered. But by definition, people learn and always want something else when they see the result. Qlik Sense has been developed with this is mind. Sense enables the IT / Reporting team to create a starting template by filling "a library of master items" with 80% of the functionality and let the business create the remaining 20% in a controlled way.

Self Service BI Process

IT / Reporting team *Build template sheets with dimension, measures and charts*

Database specialist — *Create Apps,Load and Model data*
Reporting team — *Design Charts*
Manager — *Application Validation*

Application will only be visible in personal "Own work stream"

Validate correctness and publish app into stream to make it available for the public

- App now available for the community in a specific Stream
- Sheets are marked as *"Base sheets"*

Business users *View or duplicate the sheets or charts to build new visualizations*

Consumer	Contributor	Developer
Display Only	Limited Self Service	Full Self Service
Display only • No Sheet creation	Create new sheet which is only visible for me • Use a *limited set* of dimensions • Use predefined measures	Create new Apps and sheets which are only visible for me • Use all dimensions • Use predefined measures • Create new measures • Load extra data

User can view the base sheets and the community sheets" created by the Contributor and Developer

These users can *create sheets* only visible in the users private "My sheets"
They can publish the sheet into the "community sheets" section for *collaboration* with other application users

Authorization Flow

Stream and resource access control

- Which streams are you allowed to view or modify?
- Which sheets, buttons or other resources are you allowed to view or modify. (Edit script, edit dashboard, story telling)?
- This type of authorization is managed by security rules in the management console.

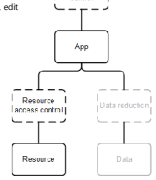

Data reduction

Answers the question: What data are you allowed to see given your userId or group?

The resource access control system in Qlik Sense is based on attributes. This means that the access is based on rules that refer to attributes connected to resources and users in Qlik Sense.

- We re-use the group membership of the source system (e.g. Active Directory)
- Each department get's its own stream (Finance, HR)
- We map the stream name to the Active Directory groups

Authorization concept

If you see the stream you may see all dashboards in it. Security is managed using one rule: Stream.name = user.group

Data Reduction

Row-level data security is accomplished by means of the Qlik Sense feature of the data model called "Section Access"

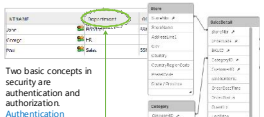

Two basic concepts in security are authentication and authorization.
Authentication answers the question "who is the user and how can the user prove it?"

Authorization answers the question "what does this specific user have access to, and what are they allowed to do".

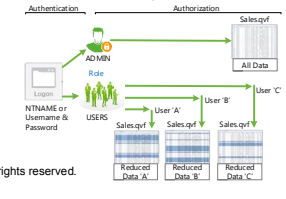

Authentication & Authorization overview

Authentication: How do we provide Single Sign On using SAML, tickets or header authentication using the Proxy API?

Authorization: What are you allowed to see/do on a specific resource?

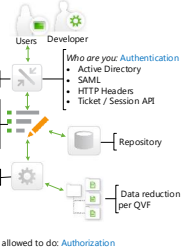

What are you allowed to do: Authorization
Re-use existing authorization or group membership definition from the following sources:

- Lightweight Directory Access Protocol (LDAP)
- Microsoft Active Directory
- Open DataBase Connectivity (ODBC)
- Qlik Sense groups
- Present a group via ticket (OEM)

Integration Overview (Embedded Analytics)

Context sharing between host website and Qlik Sense (e.g. share a session, share variables, transfer selections made) by making use of the "Session API" and "Mashup API". Integrate security to provide single sign on using web tickets, header authentication or SAML. Integrate the management/maintenance processes by using the API's

User logs into the generic portal.

Content integration using Div-tag or Iframe

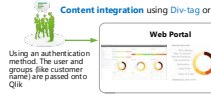

Using an authentication method. The user and groups (like customer name) are passed onto Qlik

- Iframe show a complete dashboard inside your website.
- Single integrate a single chart from an app in your site with a iFrame
- Workbench DIV tag integration, create mashups of individual charts and sheets.
- Extensions extends Sense with custom visualizations or objects

API / URL
Bi-directional communication of selections, content and user credentials.

Security Integration (SSO)

1. The user accesses Qlik Sense.
2. Qlik Sense redirects the user to the authentication module of the portal. The authentication module verifies the user's identity and credentials with an identity provider. Usually the already existing authentication mechanism of the portal which stores users/groups in a SQL table.
3. Once the credentials have been verified, a ticket is requested from Sense (QPS). Additional attributes like groups may be supplied in the request.
4. The authentication module receives a ticket.
5. The user is redirected back to the QPS with the ticket. The QPS checks that the ticket is valid and has not timed out.
6. A session is created for the user.
7. The user is now authenticated.

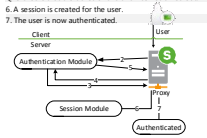

V0.18 – Date 27 Aug 2015
Created by Raymond Neeb, Qlik Presales Enterprise Architect
Contributed by Martijn Versteeg, Qlik Senior Presales Architect