

WHITE PAPER

Qlik Sense® Enterprise architecture and scalability

TABLE OF CONTENTS

Introduction	2
Associative, in-memory applications	3
User Interfaces	4
Qlik Sense Enterprise for Windows Architectural Overview	6
Multi-Node Architecture	7
Qlik Sense Enterprise with Multi-Cloud Deployment	10
Qlik Sense APIs and Development Libraries	12
Scalability at a Glance	13
Scalability Architecture	14
Scalable with Applications and Data	16
Scalable with Users	19
Scalable Delivery	20
Summary	21

PLATFORM

Qlik Sense is the Qlik® next-generation platform for modern, self-service oriented analytics. It supports the full spectrum of discovery use cases from visualization to reporting, all within a governed multi-cloud architecture. With Qlik Sense, your organization gets scalability, trust, and ongoing choice. The platform ensures your enterprise leads with data by providing broad value for all types of users with robust capabilities from unmatched Associative exploration and accelerated self-service creation to collaboration and reporting, online and offline mobility, customization and extension, data integration, and governed, multi-cloud scalability supporting the entire enterprise ecosystem. Qlik Sense runs on the patented Qlik Associative Engine, allowing all of your users of all skill levels to explore information freely without the limitations of query-based tools.

Associative, in-memory applications

Qlik couples in-memory data storage technology with an Associative Engine that lets you analyze and freely navigate data intuitively. In its second generation, the proven Qlik Data Indexing Engine allows your users to easily explore data and create visualizations based on data from multiple data sources simultaneously. These sources range from Excel® and Access® to databases such as Oracle® and SQL Server to big data sources such as Cloudera® 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 your users.

In some big data scenarios, data should remain at the source which is why Qlik uses a built-in technique called On-Demand Application Generation. Data sources can be queried based on your users' selections, yet still provide an Associative experience to your user. The Qlik Big Data Associative Index, purpose-built to expand this capability further for the biggest data sources you have, offers a high-performance, distributed multi-parallel version of the Qlik Associative Engine that can be distributed across a cluster directly at the data source.

Learn more at <https://www.qlik.com/us/bi/big-data>.

What is Associative?

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

User Interfaces

Your access to the Qlik Sense analytics environment is through a zero-footprint web browser interface. The Qlik Sense Client makes all aspects of development, drag-and-drop content creation, and consumption possible. Qlik Sense features a responsive design methodology to automatically display and resize visualizations with the appropriate layout and information to fit the device—

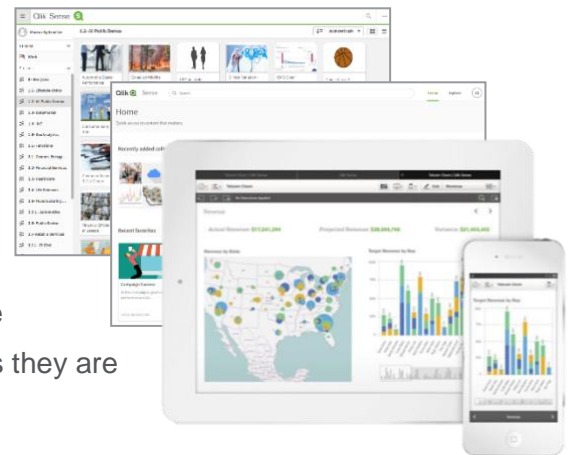


whether it is a browser on a laptop or desktop, tablet, or smartphone. Built with current standards of HTML5, CSS3, JavaScript®, and web sockets, Qlik Sense enables you to build on any platform and deploy to any platform.

The Qlik Sense Client has flexible delivery methods, depending on the use case and audience. This ensures a consistent and interactive user experience.

Hub

The Hub is the analytics portal that lets your Qlik Sense Enterprise for Windows® users access the applications they are entitled to access.



Cloud Hub

The Cloud Hub is the analytics portal that lets your Qlik Sense Enterprise multi-cloud capability users access the applications they are entitled to access.

Mobile iOS

Qlik Sense Mobile is a native iOS® application that provides online and offline analytic capabilities.

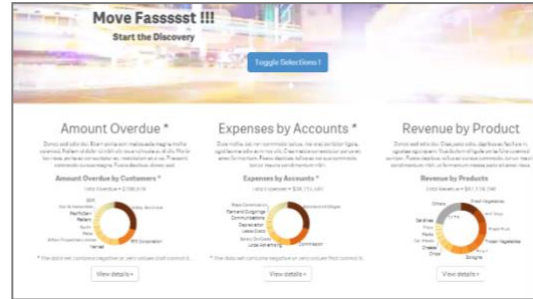
Management Console

Qlik Sense Enterprise uses the Qlik Sense Management Console. It allows your administrators to govern all aspects of the Qlik Sense platform—from data connectivity to application and task management to security administration, monitoring, and auditing.



Embedded Analytics

Qlik Sense APIs are REST and JSON based; the same APIs we used to develop Qlik Sense. You can extend Qlik Sense by adding your own visualizations and data sources. You can integrate Qlik Sense visualizations and data with other web-based applications. Qlik Sense features a built-in web-based editor for working with these APIs or use your own Integrated Development Environment (IDE).

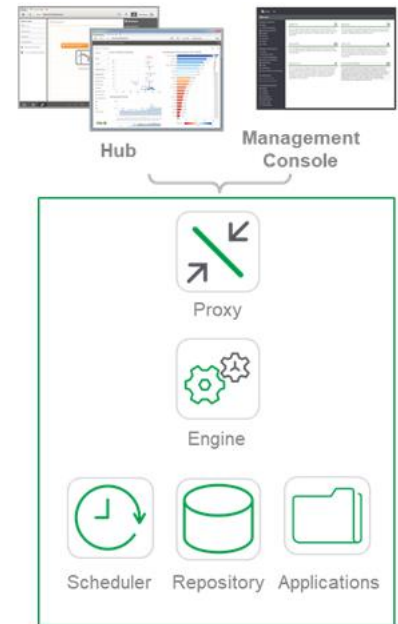


See our data sheet titled [Qlik Sense APIs](#) for more information.

Qlik Sense Enterprise for Windows Architectural Overview

Qlik Sense Enterprise for Windows is designed to support visual analytics in a scalable, secure, and governable way. You can deploy Qlik Sense on a single server and scale both vertically and horizontally to address the availability and processing requirements of your deployment, whether on-premise or in the cloud. Each node in a deployment runs some or all of the services that perform specific roles in the Qlik Sense site:

- **Qlik Sense Proxy (QPS)** is the entry point for your users and administrators via the Hub and Management Console, respectively. QPS integrates with many identity providers (e.g., Active Directory, SAML, etc.). QPS also manages your sessions, provisions your licenses, and handles load balancing to the other components.
- **Qlik Sense Engine (QES)** is the Associative, In-memory Data Indexing Engine. A RAM-based application tier, QES delivers your highly interactive, self-service visualizations, search, and calculations at runtime.
- **Qlik Sense Scheduler (QSS)** coordinates data loads. QSS supports your time-based application reloads, dependancies, and may be driven by enterprise schedulers (e.g., CA AutoSys).
- **Qlik Sense Repository (QRS)** is the platform's centralized store of configuration and management information. QRS manages your user definitions, security, and many other elements of the platform. The data persists in PostgreSQL®.
- **Qlik Sense Applications** include highly compressed data, a data model, and the presentation layer. Your applications are persistently stored on a file system and loaded into memory by the QES as your users request them.



Qlik Sense Architecture

- **Multi-Cloud Services** (not pictured) are optional services for enabling multi-cloud functionality. [See the section titled **Qlik Sense Enterprise with Multi-Cloud** for more information.]
- **Ancillary Services** (not pictured) include **Qlik Sense Printing Service (QPR)** for content exporting and **Qlik Sense Service Dispatcher (QSD)** which calls routines for data preparation, application migration, and on-demand application generation.

Multi-Node Architecture

Qlik Sense has a modular architecture designed to meet the unique requirements of your organization. While a single Qlik Sense server can often support hundreds of concurrent users, the platform's multi-node architecture makes it possible to scale even further.

See the technical brief titled [Qlik Sense Performance Benchmark](#) for more information.

Distributed Workloads

Your organization can deploy Qlik Sense components on servers in different combinations to accomplish specific functions. The table describes how you might deploy the components of Qlik Sense to allow servers to function as Consumer nodes, Developer nodes, Scheduler nodes, and the Central node. Non-Central nodes are referred to generically as rim-nodes.

Function	Description	Component service						Comment
		Proxy	Engine	Scheduler	Repository Service	Repository Database	Application Filesystem	
Consumer Node	Consumption and Self-Service Analytics for Users	X	X		X			Proxy can be separated as its own server.
Developer Node	Creation and Reloading of New Applications for Users	X	X		X			For business users. Isolated non-production environments recommended for IT-oriented authors.
Scheduler Node	Application and Data Reloads		X	X	X			Engine performs reloads invoked by scheduler.
Central Node	Master Scheduler, Application Upgrades, and Misc.	X	X	X	X	X	X	Can be co-located with other nodes.

Central Node

Every Qlik Sense environment has a Central Node that acts as the master scheduler. The Central Node coordinates application reloads, performs application upgrades when the software is updated, and serves extension objects to other nodes. By default, the Central Node also holds the Repository Database and Application Filesystem. In a multi-node environment, each Repository Service communicates with the Repository Database on the Central Node, and each Engine loads applications (.QVF) from the Application Filesystem on the Central Node. Optionally, the Repository Database and/or Application Filesystem can be moved to dedicated, resilient resources. Finally, because of its limited responsibility, the Central Node can co-locate on another node, most commonly the scheduler(s). Central Node functions can be nominated to run on one or more nodes, and in the event of a failure, one of the candidate nodes will assume Central Node responsibilities.

Single-Node Deployment



This deployment example co-locates all Qlik Sense components on a single node.

**Consumer /
Scheduler / Central
Node**

Two-Node Deployment



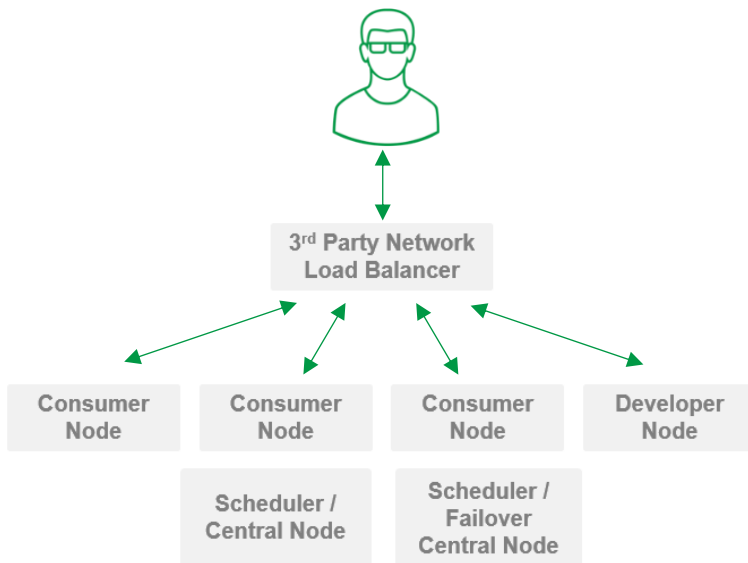
This deployment example separates the workload generated by users performing analysis from data loads managed by the scheduler.

Consumer Node

**Scheduler /
Central Node**

Scaling Horizontally – Multiple Nodes

With Qlik Sense, you can scale workloads vertically and horizontally. In a single site (environment), you can deploy multiple Consumer nodes and Developer nodes for self-service analytics as well as Scheduler nodes for data loads that 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 **Consumer and Developer Nodes**, isolate content creation from consumption or segment applications by business unit. Within **Scheduler nodes**, segment data loads by complexity or duration.

Completely configurable security rules are outlined further in the document titled

[Qlik Security Overview Whitepaper.](#)

Qlik Sense Enterprise with Multi-Cloud Deployment

Optional multi-cloud capabilities in Qlik Sense Enterprise allow your organization to broadly scale policy-driven deployments to expand the reach of analytics to new users while retaining flexibility to choose where analytic

consumption takes place. Use multi-cloud capabilities to distribute analytic applications from Qlik Sense Enterprise for Windows deployments to your own-deployed Qlik Sense Enterprise for elastic deployment environments as well as to Qlik Cloud Services, our fully managed hosted service.



Multi-Cloud Architectural Overview

Platform

The underlying architecture is microservices-based for both customer-managed Qlik Sense Enterprise for elastic deployments and our hosted Qlik Cloud Services. Either can be deployed on Linux[®] and delivered through Docker[®] and Kubernetes[®]. Our cloud-native architecture provides for elastic scale, and in the same way that we offer and manage Qlik Cloud Services, Qlik Sense Enterprise for elastic deployments can be deployed into managed container services from cloud providers or a customer-managed Kubernetes environment—enabling you to benefit from tremendous scale and unlimited potential.


Qlik follows good tenets of cloud-native design and our multi-cloud capability can be logged using standard tools such as Elasticsearch[®], Logstash[®], and Kibana[®] (i.e., an ELK stack), then monitored using standard tools such as Prometheus[®] and Grafana[®].

Learn more about microservices at [Kubernetes](#) or [Docker](#).


Components

For simplicity, microservices are categorized to denote their functions:

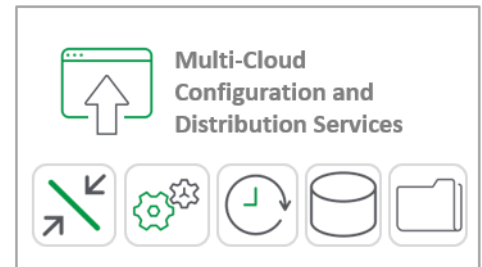
- **Cloud Hub Client** is our browser-based analytics portal serving the Qlik Sense Client for consumers in an elastic deployment.
- **Security Containers** manage authentication and entitlements to ensure your users are identified, authenticated, and only see the applications to which they are entitled.
- **Client Containers** deliver the Cloud Hub and persistent user information, preferences, and licenses.
- **Engine Containers** deliver the Qlik Sense in-memory, Associative Data Indexing Engine and supporting containers to assist with load balancing and session management. This is the same Engine in Qlik Sense Enterprise for Windows, delivered as a container.
- **Infrastructure Containers** represent basic network and datastore containers including NGINX® Ingress, Traefik, MongoDB® and Redis®.
- **Multi-Cloud Configuration and Distribution Services** are additional Windows services that become active. They are used to configure connectivity and distribute applications, entitlements, licensing, and other resources to the multi-cloud environment.



Cloud Hub



Qlik Sense Enterprise for elastic deployments



Qlik Sense Enterprise for Windows

Qlik Sense APIs and Development Libraries

Ever-changing enterprise architectures are becoming complex. Qlik Sense APIs and development libraries ensure you can integrate Qlik Sense seamlessly into your environment to provide the security, automation, and delivery capabilities your organization requires.

Qlik Sense APIs are REST and JSON based; the same APIs used to develop Qlik Sense. Use this tremendous power to build your own applications on the Qlik Sense platform. Extend Qlik Sense by adding your own visualizations and data sources. Leverage your existing security providers. Even integrate Qlik Sense with web applications or desktop applications based on the .NET framework. You can manage licensing, users, and tasks programmatically.

Qlik Sense Extensions API

- Create New Visualizations

Qlik Sense Engine / Backend API

- Communicate Directly with Associative Engine

Qlik Sense Capability API / .NET SDK

- Embed Visualizations and Functionality

Picasso.js

- Charting Library for Building Visualizations for Qlik Sense Analytics Platform

Qlik Sense QVX SDK

- Create Custom Data Connectors

Enigma.js

- Library for Interacting with the Qlik Associative Engine

Qlik Sense Repository Service API

- Management and Administration

Haylard.js

- Library for Generating Load Scripts

Qlik Sense Proxy Service API

- Authentication, Load Balancing, Session Management

Leonardo-UI.js

- Web UI Component Library with Qlik Sense Look and Feel

Qlik Sense User Directory Connector API

- Create Custom User Directory Connectors

See the data sheet titled [Qlik Sense APIs](#) and [Qlik on GitHub](#) for more information.

Scalability at a Glance

Scale matters because increasingly you need to deploy a system to larger numbers of users containing larger volumes of data and larger numbers of applications. Scalability is defined and measured along several dimensions, and is categorized into five main areas: architecture, data, applications, users, and delivery.

Qlik Sense scales and here's a look at how across the five categories. [Note: The scalability characteristics remain the same for the Qlik Associative In-Memory Engine regardless of whether it is deployed in Qlik Sense Enterprise for Windows or Qlik Sense Enterprise for elastic deployments.] Both deployment models are included in the highlights with differences noted where appropriate.

Architecture

Qlik Sense for Windows is comprised of five primary services that can be scaled vertically and horizontally to meet your demand. This simple architecture ensures predictability in performance and scaling.

Qlik Sense for elastic deployments can be deployed on Linux, and is delivered through Docker and Kubernetes. This cloud-native, microservices-based, architecture provides dynamic scaling of your Qlik deployment.

Data

Qlik Sense performance scales linearly with data volumes. Response times for analysis and ETL performance from multiple data sources are predictable as your data volumes grow.

Applications

Qlik Sense is predictable in how resource consumption scales with applications. Resource requirements of each individual application can be added to understand the requirements of a group of applications.

Qlik Sense Scalability

Architecture: Performance and capacity scale linearly with processor cores.

Data: CPU, RAM, and response times scale linearly with data.

Applications: Performance is additive, ensuring predictability.

Users: CPU, RAM, and response times scale linearly with users.

Delivery: Fewer administrative resources and self-service content creation allow Qlik Sense to scale with your organization.

Users

Qlik Sense scales linearly with the quantity of users. As your number of content creators and consumers grows, response times and impact on server resources increase predictably and linearly.

Delivery


Qlik Sense helps your organization align work with employees most suited to do it. As IT efficiently manages hardware, and secures and governs data, your users create content, self-serve, and collaborate. Fewer resources empowering many.

Scalability Architecture

Qlik Associative Engine

The unique Qlik Associative Engine couples in-memory data storage technology with a patented Associative Engine that lets you analyze and freely navigate data intuitively. This RAM-based application tier allows us to deliver highly interactive, self-service visualizations, search, and calculations at runtime.

Qlik Sense performs calculations on every click. For example, when your UI requires aggregates to show a chart object or to recalculate when selection is made, the aggregation is done at the same time. The Qlik Associative Engine is a 64-bit process, multi-threaded, and optimized to take advantage of all available processor cores. It leverages available cores linearly when performing calculations for growing data volumes and users.



Memory and Cache Management

Qlik Sense uses RAM efficiently, and in these three ways to store

- The dataset defined by the application data model
- The cached results from Engine calculations and user selections
- The session state for each user*

Qlik Sense applications persist on disk on a server. When a user requests an application, it is loaded into RAM if it has not been loaded by a previous user. Multiple users leverage this application in RAM; it is loaded once and not duplicated for multiple users accessing it. RAM is recovered when the application is no longer in use.

Memory Management

Applications

Highly compressed data; columnar storage; single copy of application data in RAM

Result Cache

Stores calculation results; automatically reused; self-managing and never out of sync with application data

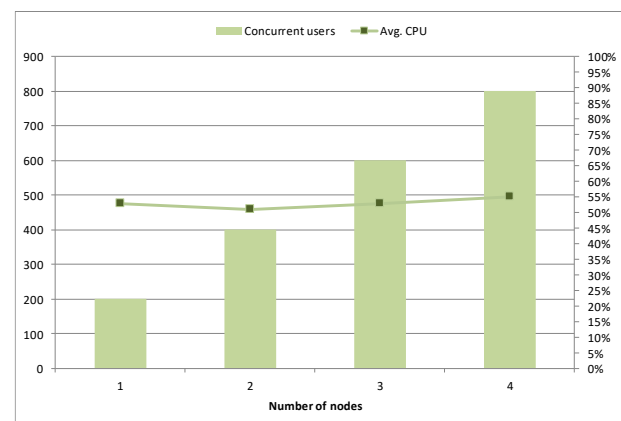
Qlik Associative Engine performs calculations in real time as users make selections within applications. RAM may be temporarily allocated to perform a calculation. Qlik stores results in a Result Cache for reuse if any user makes the same selection (provided security allows for the same data visibility). Our simple caching mechanism, which frees the CPU from performing additional, duplicate work is automatically managed by Qlik Sense. No administration required.

*Note: Each user consumes a small amount of RAM to maintain session state information.

Distributed Windows Architecture

You can deploy Qlik Sense for Windows to multiple nodes. It is supported on physical hardware, virtual platforms, and in cloud environments.

With its multi-node architecture, Qlik Sense scales linearly when additional Engines nodes are added to a deployment.



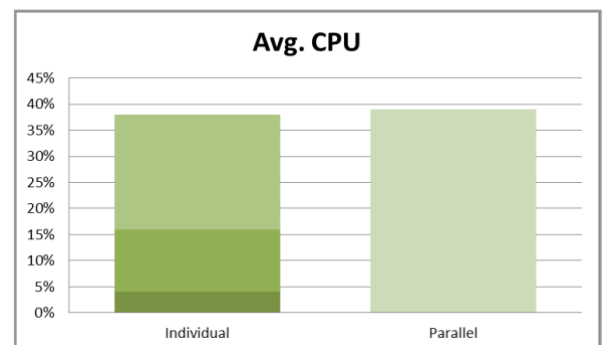
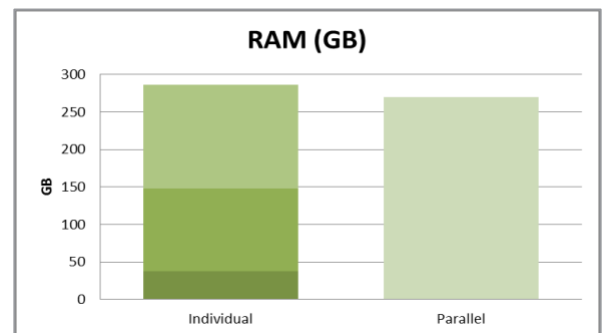
Containerized Multi-Cloud Architecture

You can deploy Qlik Sense for elastic deployments to Qlik Cloud Services, a private Kubernetes environment, or a managed Kubernetes service offered by a cloud provider. All containers can be scaled horizontally and vertically, and Engine performance scales linearly with the resources and/or number of containers allocated to it.

In either deployment model, Qlik Sense achieves scalability because it is an in-memory analytics platform that overcomes common bottlenecks prevalent in query-based tools that rely on an underlying database to perform processing. If you scale out your deployment by adding a node or increasing the number of containers and were to double processing capacity, you could handle twice the number of users (assuming a typical user produces a similar amount of load). Additionally, you can achieve even more benefits from the ability of Qlik Sense to define which applications reside on a node. Of course, many factors will contribute to scalability and overall design of your deployment. See [Planning your deployment](#) on help.qlik.com or contact us with questions.

Scalable with Applications and Data

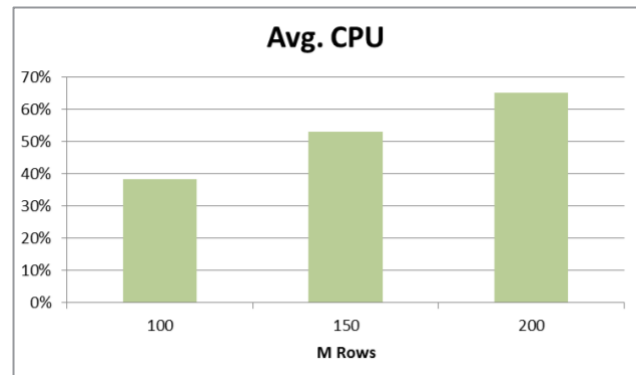
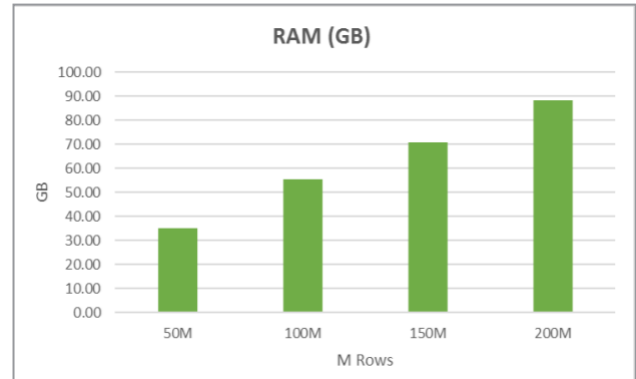
Qlik Sense deployments commonly grow to address the needs of many users across many business units. Growth requires a scalable solution for data and application management. Qlik Sense is predictable in what impact adding applications or data will have on resource consumption. Resource characteristics (CPU and RAM consumption) for single applications simply add up when deployed together, allowing you to estimate the total needs prior to publishing to a production environment. The figures illustrate how resource utilization sums predictably. For example, if you used three different applications independently on their own servers, their total RAM and CPU consumption would be equal to the same scenario if they had been served in parallel on a single server. This gives you both confidence when capacity planning and predictability while deploying.



Application Performance

As your application's source data volume grows, it requires a proportional amount of RAM to hold it. However, twice the source data does not require twice the RAM. Qlik Sense uses a highly efficient algorithm to store each unique entry once, then uses pointers to reference the relationships among data. As you add more data, you can be sure it is represented in RAM as efficiently as possible. Depending on the nature of the source data, compression results in RAM can be often up to a factor of 10.

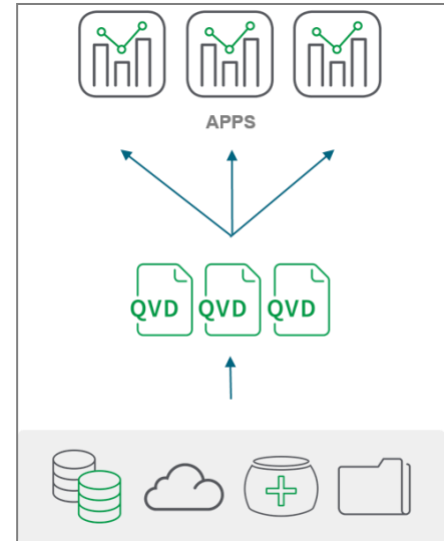
As your application's source data volume grows, it also requires a proportional number of CPU cycles to process the same calculation. Provided a result isn't already in the Result Cache, CPU cycles are used whenever a new aggregation is requested, and the amount of time needed to respond to an end-user request is a function of the CPU's ability to perform a recalculation on the data. Therefore, as you add more data to an application, the CPU will be required to perform recalculations over larger slices of data, resulting in more required CPU cycles. Fortunately, for the same reason cited above, Qlik Sense compression ensures that twice the data does not require twice the CPU power as illustrated in the graphic.



Qlik Sense performance scales uniformly with data. As you add more data to a Qlik Sense application, a corresponding addition of a combination of RAM and CPU capacity allows your end-user performance to be maintained in a predictable fashion. The linear scaling of CPU and RAM ensures you gain predictably of performance and capacity planning.

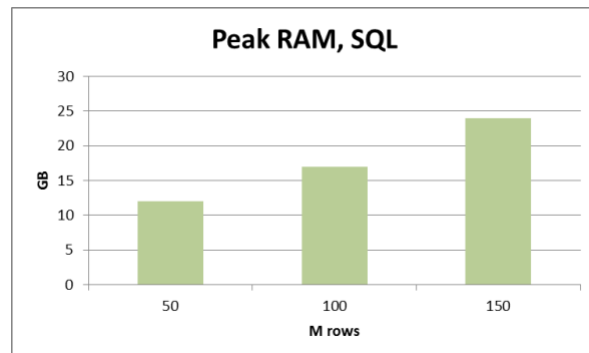

Data Layer

Qlik Sense lets you leverage an optional data layer consisting of highly compressed files called .QVD files. A benefit of this intermediate data layer is reusability and incremental loads of large data sets. Several Qlik Sense applications can use the same data in the .QVD data layer without having to individually query remote data storages, and this process, along with the full-featured ETL capability of Qlik Sense, is automated by the Qlik Sense Scheduler component.



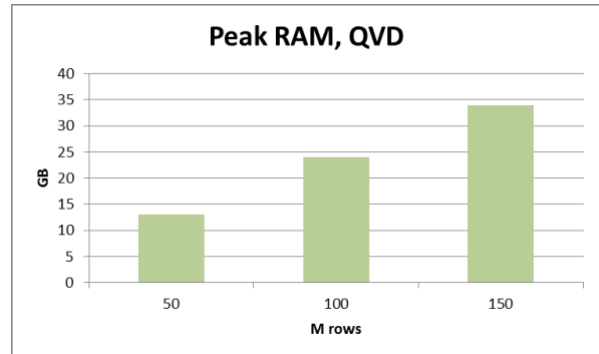
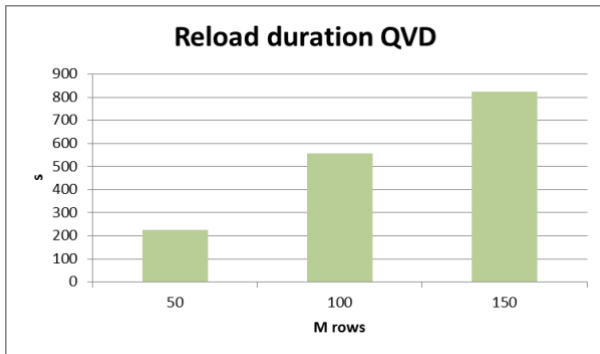
Source Data Load Performance

Qlik Sense can connect to almost any source to populate an application with data. The reload duration depends on several variables including CPU clock frequency, the data source, network speed, disk speed, types of aggregations, and calculations within load script. However, Qlik Sense scaling characteristics for reloading data show a linear trend along several dimensions as the amount of data grows. This yields predictability for reload duration. The figures illustrate measurements where different amounts of data are loaded from a database. Both duration and RAM utilization scale linearly.



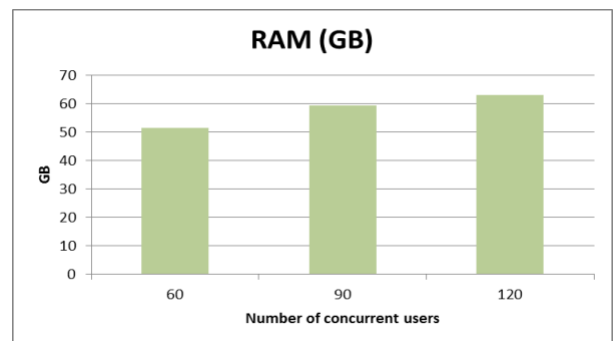
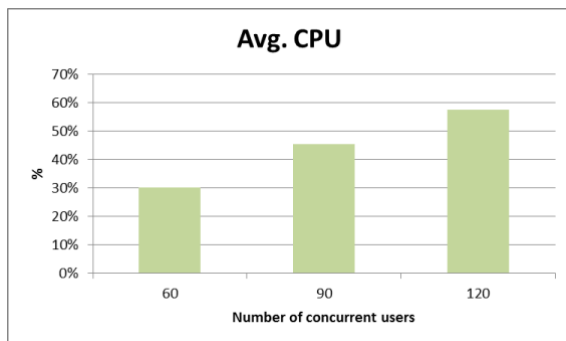
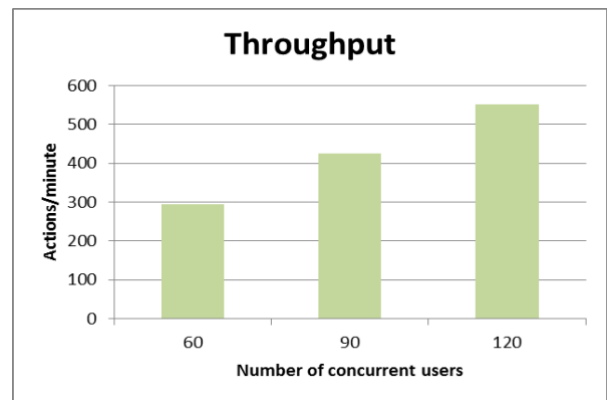
Data Layer Performance

The amount of time to load data from the Qlik Sense .QVD data layer is significantly shorter than querying the source. Loading data from the .QVD data layer also scales linearly with duration and RAM utilization.



Scalable with Users

Qlik Sense performance scales uniformly with users. As you add new users to Qlik Sense, the performance impact is predictable and addressed by a proportional addition of more CPU and RAM resources. The following graphs show that as the number of concurrent users increases, the amount of CPU and RAM consumptions increases linearly, as well.



CPU

The Qlik Sense Analytics Engine is a 64-bit, multi-threaded process optimized to take advantage of multiple processor cores when performing calculations. The instant a user makes a selection, the Qlik Sense Engine responds with a burst of intense CPU power over all CPU cores to return results as quickly as possible. When not busy, CPU utilization will approach 0%. The appropriate measure of CPU utilization and capacity then is the average CPU, and because of how the engine is designed, Qlik Sense can scale CPU utilization linearly as you add users to the platform.

RAM


The RAM chart illustrates principles highlighted in the section titled Memory and Cache Management. As the number of concurrent users doubles, the RAM footprint increases linearly, but the actual amount of RAM used does not. Application data is loaded a single time and not duplicated for multiple users accessing and analyzing it. Note RAM utilization actually slows as the number of users increases. This is an effect of increasing hits on the global cache.

Scalable Delivery

Qlik Sense helps your organization align work with the people most suited to do it. Your IT professionals can efficiently manage hardware by deploying Qlik Sense in virtual, cloud, and physical environments. The platform can be reconfigured, scaled up, out, or scaled down as needed. Your IT professionals can control access to data sources through centrally managed data connections, and delegate administrative capabilities as appropriate to line-of-business administrators at different levels.

Within your user community, content creators can combine and prepare data, build and use libraries of dimension and metrics, drag and drop visualizations, create stories, and publish content.

With Qlik Sense, your organization has the freedom to lead with data, discovering insights that disrupt industries by deploying self-service at scale.



Summary

Qlik Sense offers consistent and predictable performance through an architecture that gives your organization flexibility to deploy where you need and to scale horizontally and vertically as well as linearly with data, applications, and users. Using Qlik Sense, your IT professionals effectively maintain end-user performance while having the confidence to plan capacity for future, expected system usage.

To learn more, visit qlik.com.



About Qlik

Qlik is on a mission to create a data-literate world, where everyone can use data to solve their most challenging problems. Only Qlik's end-to-end data management and analytics platform brings together all of an organization's data from any source, enabling people at any skill level to use their curiosity to uncover new insights. Companies use Qlik 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 48,000 customers around the world.

qlik.com

© 2018 QlikTech International AB. All rights reserved. Qlik®, Qlik Sense®, QlikView®, QlikTech®, Qlik Cloud®, Qlik DataMarket®, Qlik Analytics Platform®, Qlik NPrinting®, Qlik Connectors®, Qlik GeoAnalytics®, Qlik Core®, Associative Difference®, Lead with Data™, Qlik Data Catalyst™, Qlik Associative Big Data Index™ and the QlikTech logos are trademarks of QlikTech International AB that have been registered in one or more countries. Other marks and logos mentioned herein are trademarks or registered trademarks of their respective owners. QSEARCH120418_AM