

# Performance Benchmarks

For Qlik Sense Enterprise on Windows

## TABLE OF CONTENTS

---

Summary	2
Introduction	2
Quick Facts	3
Qlik Sense Scalability Tool	4
Applications	5
Overall Score	6
Conclusion	8

## Summary

---

- These performance benchmarks are applicable to Qlik Sense Enterprise and Qlik Analytic Platform when deployed on Windows
- The published results should only be used as general guidance when planning deployments. Your results may differ due to other factors and variables within your own environment such as; machine specification, size and number of apps, number of users, etc.
- The Qlik Capacity Benchmark gives a clear indication around the data volumes and concurrent users that can be handled when a server is taken to saturation, but also show metrics when a server is not saturated, as well.
- A complete and transparent view of the performance of Qlik Sense Enterprise on Windows is provided through fundamental and critical metrics of CPU, RAM, and response times.

## Introduction

---

This technical brief contains performance benchmarks for Qlik Sense based on a testing methodology called the Qlik Capacity Benchmark. A series of tests are conducted by varying data volumes, users, and applications in a given environment. The results are an exhaustive set of permutations recorded in a matrix of CPU utilization, RAM utilization, and response times. This approach differs from many other scalability tests as it shows reported metrics of when a server is saturated, as well as when it is only partially utilized. This methodology provides transparency to the testing process and resulting metrics. While Qlik applications do vary in size and complexity, the methodology provides a comprehensive set of data with which you can judge scalability and plan for your deployments.

For more information about Qlik Sense Enterprise on Windows, please reference the [Qlik Sense Enterprise - Architectural and Scalability Technical Whitepaper](#).

## Quick Facts

This benchmark demonstrates capacity and performance with variations of concurrent users, data volumes, and applications complexities listed below.

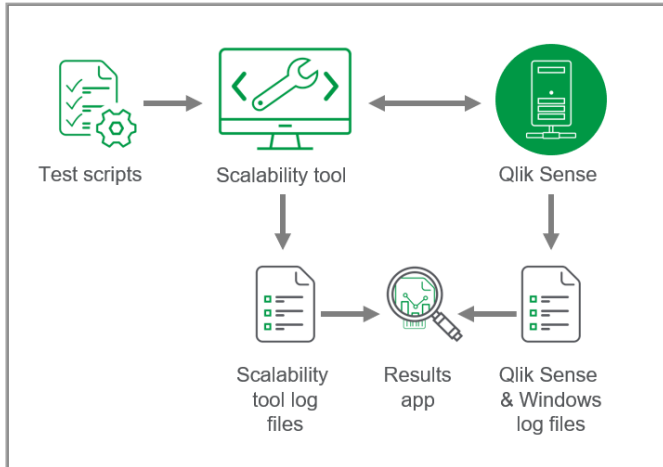
Scenario	Software	Instance type	RAM	Engine Cores	Concurrent Users	Data Volumes	Application Complexity
1	Qlik Sense Enterprise June2019	r5.8xlarge	256	32 vCPU	10 - 500	10 – 500M rows	moderate, complex
2	Qlik Sense Enterprise June2019	r5.2xlarge	61	8 vCPU	10 - 200	10 – 200M rows	moderate, complex
3	Qlik Sense Enterprise June2019	r5.2xlarge	61	8 vCPU	1000	1M rows	Simple

Notes:

- EC2 load cannon details: m4.4xlarge (64 vCPU/64GB RAM)
- These performance benchmarks are applicable to Qlik Sense Enterprise and Qlik Analytic Platform when deployed on Windows.

# Qlik Sense Scalability Tool

Qlik Sense was tested with a freely available load testing tool called Qlik Sense Scalability Tool. It is found on [Qlik Community](#). This load testing client can simulate virtual users against customer Qlik Sense applications.



Using the Qlik Sense Scalability Tool, test scripts simulate virtual users and are executed against Qlik Sense applications. Upon completion, performance metrics from Qlik Sense, Windows, and the Scalability Tool are collected into a Qlik application for analysis. All virtual users were simulated to be highly interactive with the application. In all scenarios, virtual users interacted with charts and filter panes, navigated

among tabs, and performed actions within applications. This provided a realistic view of how Qlik Sense handles a given user load. Virtual users were simulated with 20 to 30-second think times and made randomized selections throughout the tests, rather than the same selection, to minimize caching that might under report utilization averages. Each test ran for at least one hour.

## Control Variables

The control variables varied throughout a series of tests:

- **Applications:** moderate, complex
- **Concurrent Users:** 10, 25, 50, 75, 100, 200, 500, 750
- **Data Volumes (Millions):** 10, 50, 100, 200, 500

## Metrics

The metrics captured during each test:

**Average Engine CPU Utilization:** 0 to 100%

**Max RAM Utilization:** 0 GB to Max GB of Server

**Average User Response Time:** 0 sec to 5(+) sec

Test #	Application	Concurrent Users	Data Volume
1	moderate	10	10
2	moderate	50	10
...	...	...	...
6	moderate	10	50
...	...	...	...
12	complex	10	10
...	...	...	...
56	complex	50	500

Results

Test #	CPU %	RAM (GB)	Response Time
1	4	4	0.5
2	8	5	0.5
...	...	...	...
6	12	12	0.5
...	...	...	...
12	6	6	0.8
...	...	...	...
56	65	200	1

Example Capacity Benchmark Matrix

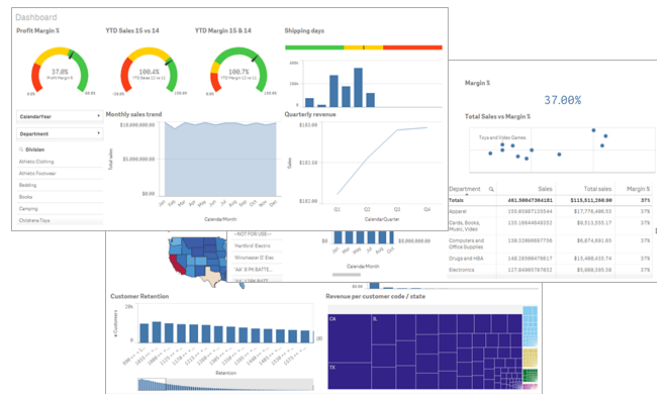
# Applications

The complexity of an application influences how many concurrent users and how much underlying data it can support. The Qlik Capacity Benchmark tests account for this variation by testing applications with different presentation layers, calculations, and test scripts.

The applications used in the tests are benchmark apps designed based on our experience to reflect common scenarios that you would typically see within your organization.

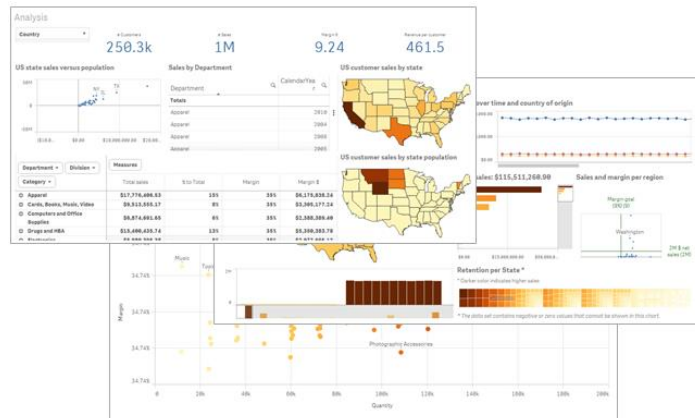
## Sales Dashboard (Moderate)

The Sales Dashboard application shows data in aggregate via many graphical objects. In addition to gauges and trends, it allows for more complex analysis including mapping, scatter charts, and set analysis. The client load script simulates a use case where users research data at an aggregate level, drill through many contexts (customer, profitability), and interact with charts.



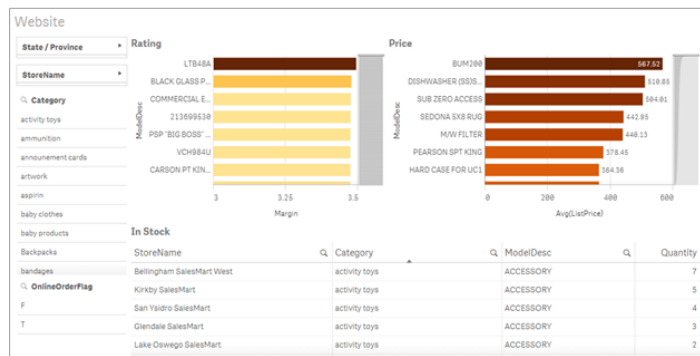
## Sales Analysis (Complex)

The Sales Analysis application allows for analytics including many calculations per sheet, dense scatter plots, and some detail data throughout the application. The client load script simulated a use case where users perform complex analytics, including set analysis, and product ranking.



## Retail Website (Simple)

The Retail Website application shows a simple scenario where the user researches product location, availability, and pricing on a limited dataset of 1 million rows.



# Overall Score

The results from the Capacity Benchmark are categorized according to the thresholds defined in the table to the right. The primary metrics of average engine CPU utilization, maximum RAM utilization, and average response times are scored in this way to provide visual feedback about the overall performance of the server in each scenario. The tests are unbounded and are not intended to show maximum performance. The scores also give an indication of the overall remaining capacity of the server in each configuration. Finally, the scores are rolled into an overall score for the server.

Score	CPU	RAM	Response Time
●	< 60 %	< 70 %	< 1 s
●	< 80 %	< 90 %	< 3 s
●	> 80%	> 90 %	> 3 s

In the results below, for example, a green mark indicates that the test completed with less than 60% average CPU utilization, 70% RAM utilization, and less than one second response time. A yellow mark indicates one or more metrics entered the yellow range, and a red marking indicates one or more metrics entered the red range.

In addition, this document doesn't comment on the total *addressable* data per server. With application architectures, such as On-Demand-App-Generation (ODAG) and Direct Discovery, the total addressable amount of source data could be much larger. For more information please see the [Qlik and Big Data whitepaper](#) and also visit [Qlik Community](#).

*Note there is nothing inherently wrong with a server running more than 60% CPU or 70% RAM utilization; we do this to show a realistic viewpoint of the remaining server capacity under a given load from a sizing and capacity planning standpoint. As shown below, tests with 'yellow' CPU or RAM utilization still yield acceptable response times but may be limited in additional capacity. 'Red' indicates system saturation.*

## Scenario #1 - Capacity Benchmark Summary – 32 cores

AWS r5.8xlarge, 32vCPU, 256 GB RAM

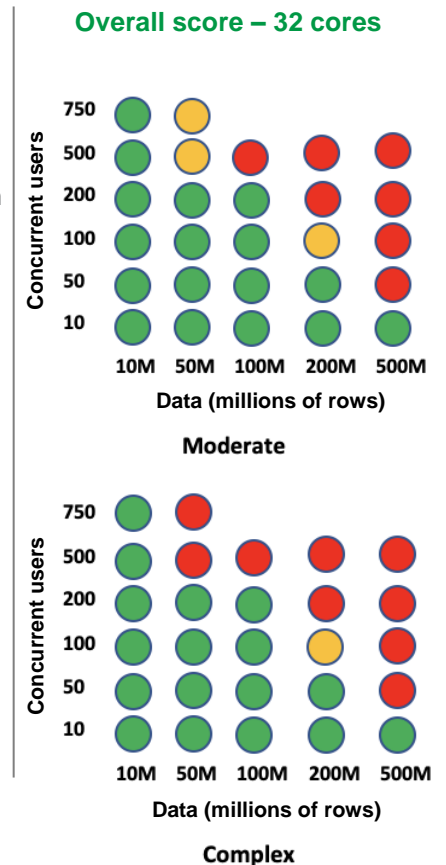
Qlik Sense was tested with 54 performance tests split equally between moderate and complex applications. The overall scores are displayed to the right where each point represents at least an hour-long test.

### Results

Qlik Sense was able to reach 750 concurrent users (7500+ total users) on a 50 million row data set and 10 concurrent users (100+ total users) on a 500 million row data set.

Data High Water Mark					
Application Complexity	Data Volume (Millions of Rows)	Concurrent Users	Avg Response Time (Seconds)	Average CPU %	Max Ram (GB)
moderate	500	10	0.7	17.28	53.09
complex	500	10	0.7	24	89.51

User High Water Mark					
Application Complexity	Data Volume (Millions of Rows)	Concurrent Users	Avg Response Time (Seconds)	Average CPU %	Max Ram (GB)
moderate	50	750	1.6	26.7	46.5
complex	10	750	0.1	13.8	35.7



## Scenario #2 - Capacity Benchmark Summary – 8 cores

AWS r5.2xlarge, 8 vCPU, 61 GB RAM

Qlik Sense was tested with 51 performance tests for moderate and 38 performance tests for complex applications. The overall scores are displayed below, where each point represents at least an hour-long test.


### Results

Qlik Sense was able to reach 400-500 concurrent users (4000-5000+ total users) on a 10 million row data set and 25-125 concurrent users (250-1250+ total users) on a 200 million row data set.

Data High Water Mark					
Application Complexity	Data Volume (Millions of Rows)	Concurrent Users	Avg Response Time (Seconds)	Average CPU %	Max Ram (GB)
moderate	200	125	2.5	75.06	33.6
complex	200	25	2.5	68.2	41.5

User High Water Mark					
Application Complexity	Data Volume (Millions of Rows)	Concurrent Users	Avg Response Time (Seconds)	Average CPU %	Max Ram (GB)
moderate	50	500	2.5	54.6	33.2
complex	10	400	1.8	53.4	28.4

### Overall score – 8 cores





## Scenario #3 - Retail Website Application – 8 cores

AWS r5.2xlarge, 8 vCPU, 61 GB RAM

The Retail Website application shows a simple scenario where the user researches product location, availability, and pricing on a limited dataset of 1 million rows. This is intended to illustrate the scale at which Qlik operates scenarios where the data set is very small.



Test Input
1 million rows of data
1000 concurrent users

Metric	Value
Engine CPU %	14.57%
Average Response time	1.044 seconds

## Conclusion

The Qlik Capacity Benchmark tests are different from many other scalability tests. Not only is a clear indication given around the data volumes and concurrent users that Qlik Sense Enterprise, when deployed on Windows, can handle when a server is taken to saturation but also show metrics when a server is not saturated, as well. These fundamental and critical metrics of CPU, RAM, and response times provide a complete and transparent view of the performance of Qlik Sense.



## 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](https://www.qlik.com)

© 2019 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™, Qlik Insight Bot™ and the QlikTech logos are trademarks of QlikTech International AB that, where indicated by an "®", have been registered in one or more countries. "Attunity" and the Attunity logo are trademarks of Attunity Ltd. Other marks and logos mentioned herein are trademarks or registered trademarks of their respective owners.