



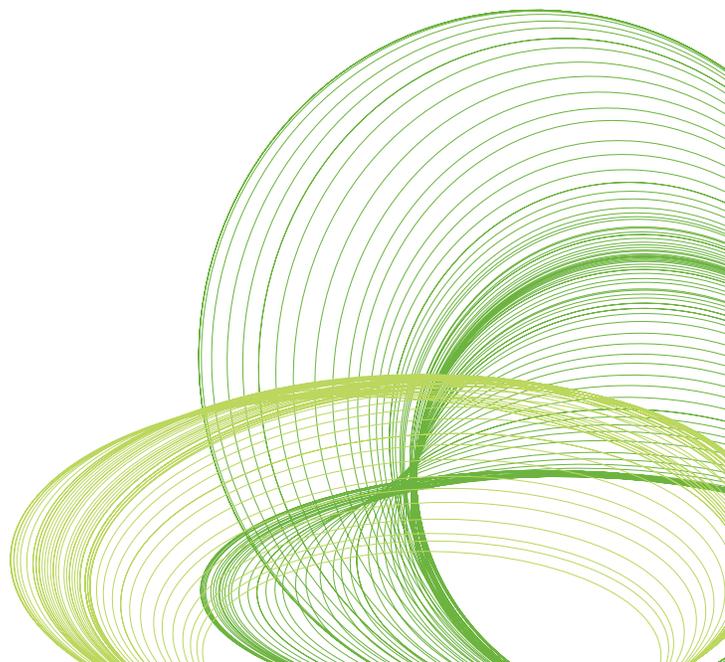
# QLIKVIEW DEPLOYMENT AT A LARGE MULTINATIONAL BANK

---

QlikView Technical Case Study Series — Scalability

September, 2011

[qlikview.com](http://qlikview.com)



## Introduction

---

This document describes the QlikView deployment for a multinational bank at a technical level. This deployment serves the business discovery needs of six thousand users in eight regions analyzing 1.4 TB of financial data for private individuals, corporations and organizations. This paper illustrates one of the real-life examples of QlikView deployment with large numbers of users located in different geographies, containing large volumes of data to demonstrate how QlikView scales in these types of deployment scenarios. Please note that the technical specifications of a QlikView deployment depends on a range of considerations such as the size and the nature of the source data, the number of concurrent users, GUI design and overall applications data design. The technical deployment decisions explained in this paper are dependent on the considerations specific to this customer. It is a referral paper and should not be used as a deployment guideline.

The first section of the paper, Technical Solution, provides an understanding of the QlikView architecture, hardware specifications, details on the data source and size, number of users and license types.

The second section of the paper, Performance Results, illustrates how QlikView scales predictably and examines end user response time.

Before reading this paper, it is recommended that you review the QlikView Architecture and System Resource Usage Technical Brief in order to get a fundamental understanding of the various QlikView components and how they utilize various hardware resources such as RAM and CPU\*.

This technical deployment case study is a companion piece to the QlikView Scalability Overview Technology White Paper as it provides a fundamental grounding from which a better understanding of how QlikView scales can be gathered. It is recommended that you download and read the Scalability Overview Technology White Paper after reading this technical deployment case study\*\*.

\* [www.qlikview.com/.../DS-Technical-Brief-QlikView-Architecture-and-System-Resource-Usage-EN.ashx](http://www.qlikview.com/.../DS-Technical-Brief-QlikView-Architecture-and-System-Resource-Usage-EN.ashx)

\*\* <http://www.qlikview.com/us/explore/resources/whitepapers/qlikview-scalability-overview>

## LARGE SCALE QLIKVIEW DEPLOYMENT AT A MULTINATIONAL BANK

- 6150 business users located in 8 regions
- 1600 concurrent users
- 1.4 TB of financial data
- 3088 clicks per minute
- Average response time less than 3 seconds

## QLIKVIEW PRODUCTS COMPONENTS

- 3 clustered QlikView servers
- 1 standalone QlikView server for heavy analytics
- 2 clustered QlikView publishers
- QlikView developer and zero-footprint Ajax clients

## Executive Summary

---

The ability to scale and support thousands of users has become an increasingly important requirement for BI platforms. This document covers the scalability approaches a large multinational bank took for their QlikView deployment to support six thousand users' business discovery needs, analyzing 1.4 TB of financial data.

The business users are dispersed in eight different regions. With QlikView's uniformly scaling capability, the business discovery environment was able to support on average 1600 concurrent users with average response times under three seconds.

## Technical Solution

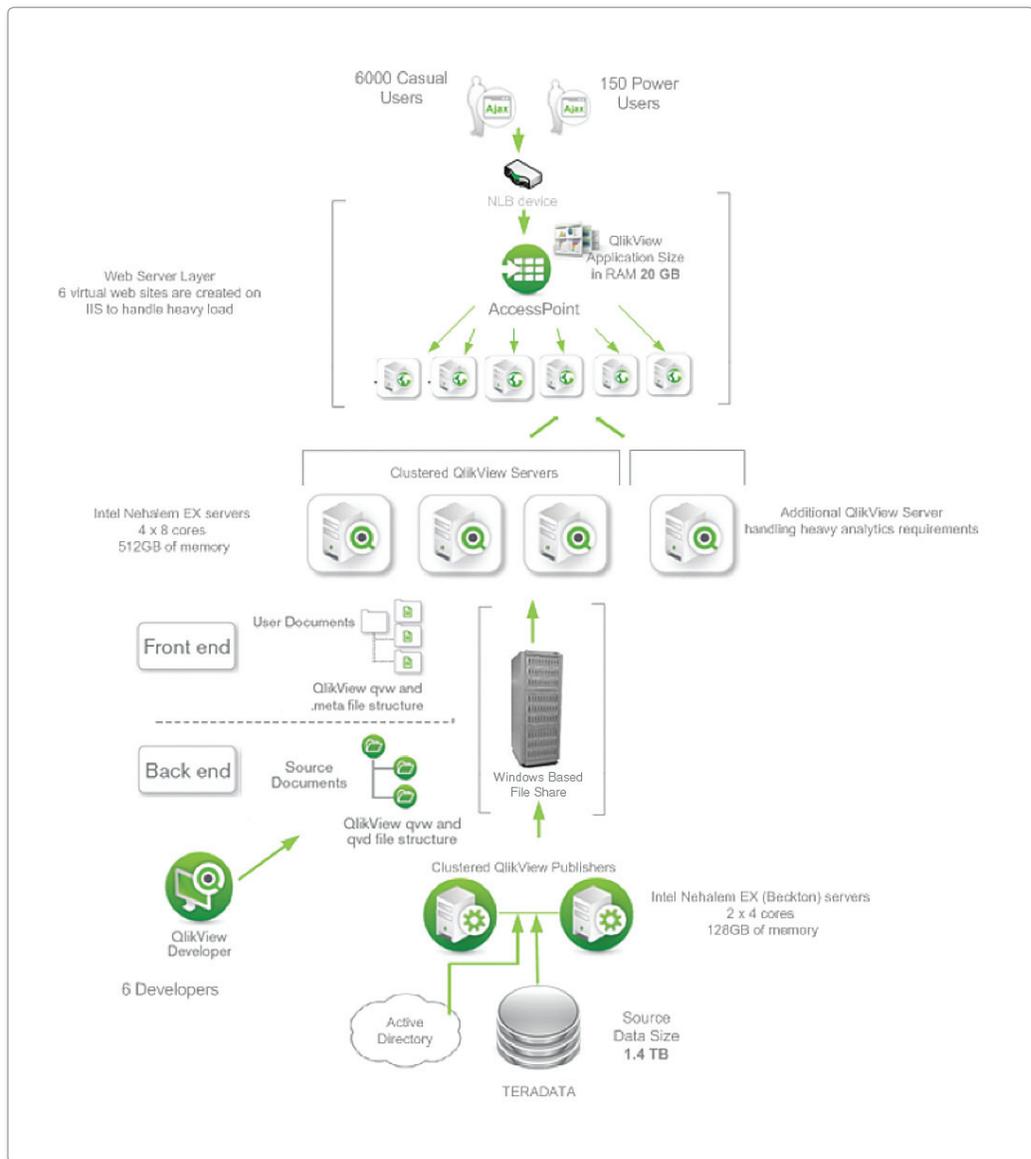
---

### ARCHITECTURE OVERVIEW

Figure 1 shows the architecture of the QlikView production environment. Three main infrastructure components are used in this deployment.

- **QlikView Server (QVS)** handling the communication between clients and the QlikView applications. There are four servers used to host the QlikView Server, three servers are clustered and one server dedicated to power users to handle heavy analytics requirements. All of the servers reside behind a network load balancer which directs the traffic. All four servers share the same storage area (Windows Based File Share) to access the QlikView applications.
- **QlikView Publisher** is used to load the data from Teradata. Initial size of the source data is 1.4 TB. Two servers are clustered with QlikView Distribution Service to handle the data refreshes.
- **QlikView Clients:** QlikView Developer is used to develop the data extraction and transformation layer as well as to develop the graphical user interface. As the business user access required no installations, QlikView's AJAX zero-footprint browser client is used for business users accessing QlikView applications on the QlikView Server.

**Figure 1: QlikView Production Environment Architecture**



© 2011 QlikTech

This deployment illustrates one of the QlikView best practices in terms of separating the front end and the back end of the system.

- The back end of the system holds QlikView source documents (QVW and QVD files), created using the QlikView Developer. QlikView Publisher is in charge of data reloads and application distribution. The Windows file system is used for QlikView source document authorization.
- The front end is where end users use the Ajax client to interact with QlikView documents that they are authorized to see via the QlikView Server. It contains the QlikView user documents that have been created via the QlikView Publisher on the back end.

## **HARDWARE SPECIFICATIONS**

### **QlikView Servers (front end)**

- 4 Intel Nehalem EX servers
- 4 x 8 cores
- 512 GB of memory

### **Publisher Servers (back end)**

- 2 Intel Nehalem EX servers
- 2 x 4 cores
- 128 GB of memory

## **CLUSTERING DETAILS**

- QlikView Server cluster to load balance based on RAM and CPU availability
- Web server clustering to cope with heavy load on the IIS web server
- Cisco NLB

Please note that the technical specifications of this QlikView deployment depend on a range of considerations specific to this customer. The size and the nature of the source data, the future growth rate of concurrent users and the GUI design impact the technical details such as server and web server specifications. Please refer to QlikView Scalability Overview Technology white paper to understand the driving dimensions that influence the technical details of QlikView deployments.

## **HARDWARE SPECIFICATIONS**

Four QlikView Servers (QVS) are used at the front end of the deployment. These are Intel Nehalem EX servers with 4 x 8 cores and 512GB of memory.

There are two publisher servers used at the back end. These are Intel Nehalem EX servers with 2 X 4 cores and 128GB of RAM.

One of the main considerations on determining the server size was to retain performance levels as the data, the number of users and QlikView applications grow in the future.

## **CLUSTERING DETAILS**

Web server clustering is used to cope with heavy load on the IIS web server. In order to handle six thousand users and the future user base growth, six virtual web sites are created. Each site uses its own application pool thereby spreading the load.

Six virtual web sites run behind a Cisco network load-balancer to handle heavy load on QlikView Accesspoint. The user sessions are sent to each machine based on the machine's availability and is accomplished by round robin.

QlikView Server clustering is used to support load sharing of QlikView applications across the three servers. With QlikView clustering, the business user requesting the QlikView application is directed to that server which has the most RAM and CPU available.

## NUMBER OF USERS

- 6000 business users
- 150 power users
- 6 developers

## DATA SIZE

- 1.4 TB of Teradata
- QlikView application size in RAM 20 GB

## NUMBER OF USERS AND LICENSE TYPES

The deployment supports six thousand casual users, one hundred and fifty power users and six developers. The Named User CAL is the only client license used.

## QLIKVIEW APPLICATION DETAILS

The initial source data size of the QlikView application is 1.4 TB. With QlikView's sophisticated data compression algorithm (as outlined in the QlikView Architecture and System Resource Usage Tech Brief), the size of the QlikView application in RAM is only 20 GB. The total number of records in the QlikView application is approximately 500 million. Dynamic data reduction is used to dynamically reduce the number of rows for each user opening the QlikView application based on their region and branch information.

## SECURITY DETAILS

This deployment uses Active Directory both for authentication and authorization. Authentication to the QlikView Server is based on the Integrated Windows Authentication (IWA) and the logged-in user's identity is communicated to QlikView Server using NTLM. This solution uses single sign-on capabilities right out of the box.

Once the user has been authenticated, the QlikView Server controls access to a given QlikView document by determining if the authenticated user has NTFS privileges to the underlying QlikView Document file (.QVW file). In this deployment the default document authorization model, NTFS, is used as all users and groups are identified in Active Directory.

In addition to these, a QlikView technology for data-level security called Section Access is used to create row level security on the QlikView application. Section access creates an internal security table on the QlikView document which stores the information on which user can see which branch of data. QlikView uses this information to dynamically reduce the data, so the user only sees his own branch's data.

## STORAGE DETAILS

Network attached storage (Windows Based File Share) is used as the storage device for QlikView documents (QVWs and QVDs). It has two different shares and file structures, separating the front end and the back end systems.

## PERFORMANCE SIMULATION RESULTS

- 500 million rows in QlikView application with 1.4 TB total source data size
- 1600 concurrent users
- 3088 clicks per minute
- Average response time less than 3 seconds

## Performance Results

---

This section summarizes some of the performance results seen while testing the deployment. An automated performance and testing product is used to put the QlikView application through the rigors of real-life user loads, while collecting performance information from the QlikView Server. The results were then analyzed in detail to understand the average response time of the QlikView application.

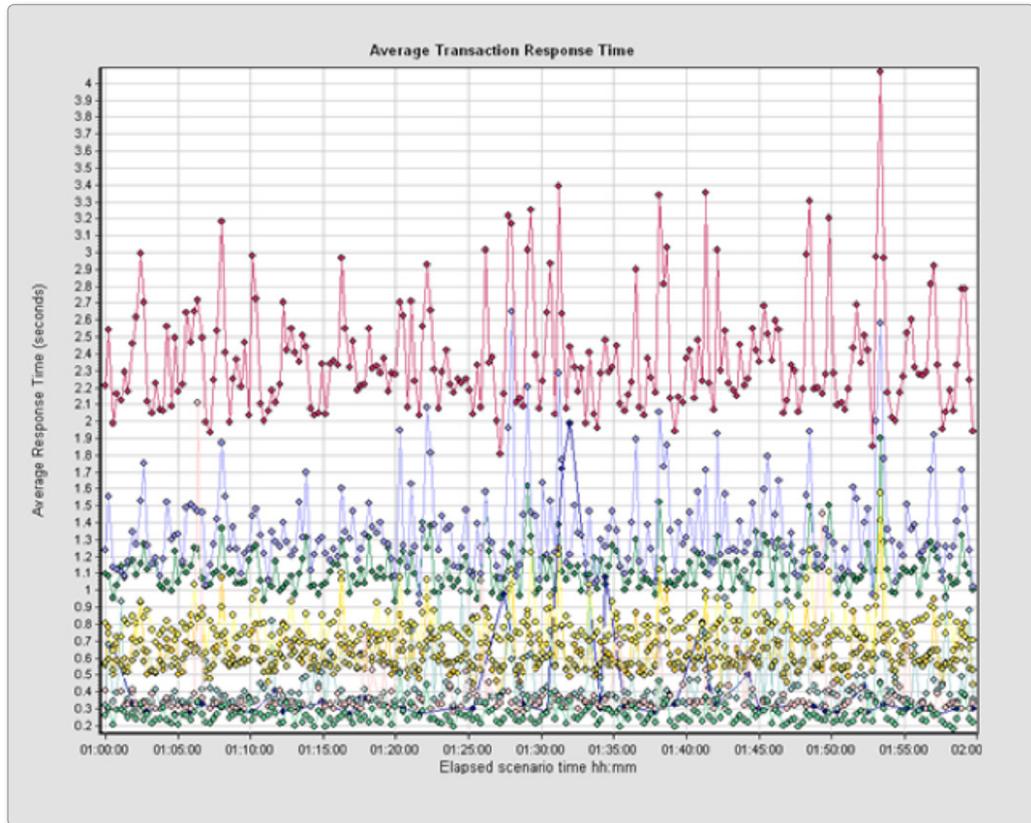
The QlikView application used in the simulation was 20 GB when opened in RAM. The test emulated 1600 concurrent users applying different actions on the application at the same time. The test was setup to emulate 3088 clicks per minute on the application. Below are some of the different actions simulated with each click.

- Open QlikView application
- Select tabs on the QlikView application
- Select values from different list boxes on a tab
- Apply bookmarks
- Clear all selections

Please note that QlikView has unique in-memory technology that allows storing all of the data at granular level in memory and all of the aggregations and calculations are done on the fly on the charts for the specific selections done by business users. Please refer to QlikView Architecture and System Resource Usage Technical Brief in order to get a fundamental understanding of the QlikView's unique associative in- memory technology.

Figure 2 summarizes the average response times in seconds for different actions simulated on the application with minimum, maximum and standard deviation information. The average response time for all of the actions was under 3 seconds.

**Figure 2: Performance simulation results in seconds for emulating 1600 concurrent users with 3088 clicks per minute**



| Color | Action                    | Minimum | Average | Maximun | Std. Deviation |
|-------|---------------------------|---------|---------|---------|----------------|
| ●     | Open Qlikview Application | 0.703   | 1.129   | 10.11   | 0.596          |
| ●     | Apply Bookmark            | 1.453   | 2.391   | 9.703   | 0.858          |
| ●     | Apply Bookmark            | 0.234   | 0.65    | 6.813   | 0.569          |
| ●     | Selections on List Box    | 0.266   | 0.43    | 1.984   | 0.361          |
| ●     | Selections on List Box    | 0.266   | 0.444   | 2.656   | 0.303          |
| ●     | Clear All                 | 0.375   | 0.596   | 4.922   | 0.297          |
| ●     | Selections on List Box    | 0.266   | 0.378   | 2.109   | 0.191          |
| ●     | Tab Change                | 0.625   | 1.364   | 8.547   | 0.747          |
| ●     | Back Selection            | 0.409   | 0.806   | 7.656   | 0.533          |
| ●     | Expand Charts             | 0.125   | 0.263   | 4.578   | 0.266          |

## Conclusion

---

This paper has outlined the technical QlikView deployment details taken at a multinational bank to handle the needs of a large business user community. This deployment is just one example of many large QlikView deployments proving QlikView's scalability to handle many thousands of concurrent users and large quantities of data whilst achieving high end-user performance.

This large multinational bank implemented QlikView to support six thousands business users' business discovery needs for analyzing financial data. The business users are dispersed in eight different regions. Their QlikView environment includes 3 clustered QlikView Servers, one standalone Server to handle heavy analytics requirements, and 2 clustered Publisher servers handling 1.4 TB of data. The performance tests conducted on QlikView Servers confirm that QlikView maintains end-user performance at the expected high levels with the average response time per click under three seconds.

## References

---

QlikView Architecture and System Resource Usage Technical Brief

[www.qlikview.com/.../DS-Technical-Brief-QlikView-Architecture-and-System-Resource-Usage-EN.ashx](http://www.qlikview.com/.../DS-Technical-Brief-QlikView-Architecture-and-System-Resource-Usage-EN.ashx)

QlikView Scalability Overview Technology White Paper

<http://www.qlikview.com/us/explore/resources/whitepapers/qlikview-scalability-overview>