White Paper
April 2017

Qlik NPrinting™ 17.x: Architecture and Scalability

Contents

About Qlik NPrinting

Qlik NPrinting 17 is Qlik's next-generation reporting and distribution solution for Qlik Sense® and QlikView®, helping organizations spread knowledge and insight by creating and distributing great looking reports. Qlik NPrinting (or "NP") offers reports in widely adopted, portable formats such as PDF, HTML, Excel, Word and PowerPoint. Through scheduled distribution and a new online subscription hub, Qlik NPrinting ensures that the right reports get to the right people, how and when they need them. Qlik NPrinting's new multi-threaded engine can be configured across multiple servers to increase the scalability of your Qlik NPrinting investment.

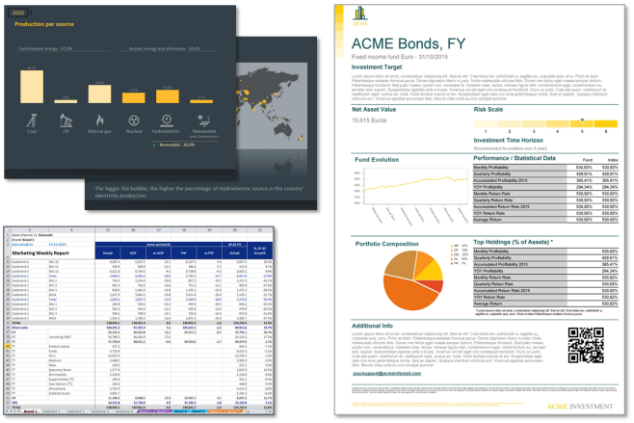

Architecture

Qlik NPrinting is a server based component that can be added to an existing QlikView and/or Qlik Sense environment. Qlik NPrinting is not a stand alone solution, it can only be used in conjunction with QlikView or Qlik Sense.

The Qlik NPrinting server is responsible for sourcing data securely from QlikView or Qlik Sense and producing and distributing reports in static format. It consists of a central server component with one or more engine components that can be distributed as a cluster over multiple servers. It also includes a desktop designer, as well as the web console and an end user web based NewsStand portal. A web based API is also available which allows for automation and management of the platform as a whole.

The server-side components and client interfaces are described below.

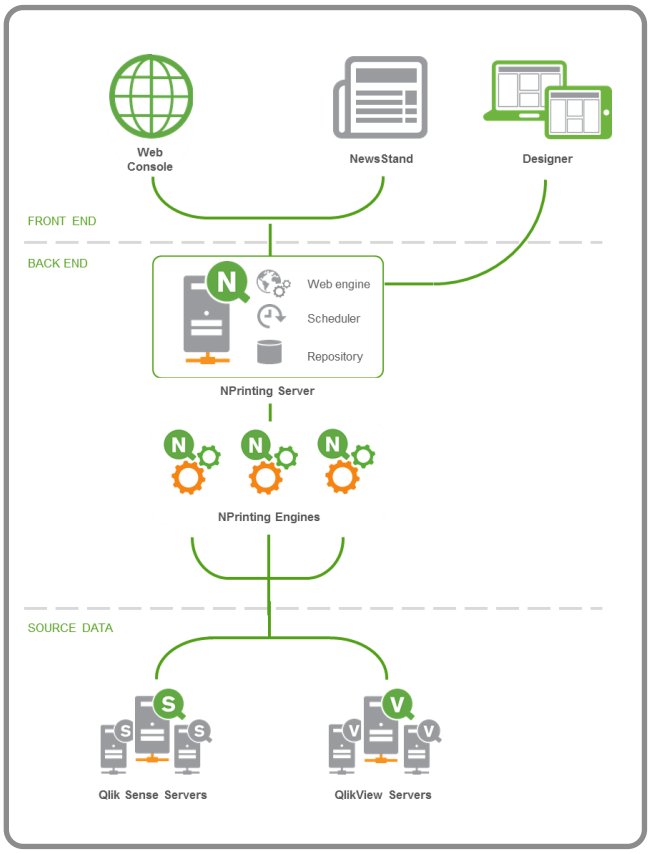

Front End

The Web Console and NewsStand user interfaces are web-based and can be accessed using a modern web browser.

Web Console – The Web Console is an administrative console for administrator and developer use only.

NewsStand - The NewsStand web portal allows users to check, download, and subscribe to reports for which they are authorized.

Qlik NPrinting Designer - The Designer is a desktop application that allows developers to create and manage report templates. It is launched from the web console. The designer interface is a local client.

Back End

Qlik NPrinting Server – Only a single NP Server is required in a Qlik NPrinting site, which is comprised of the Scheduler, Repository, Web engine, and Messaging Service

- **Scheduler** - The scheduler service performs the following tasks:
  - distributes jobs among available engines
  - manages job prioritization
  - delivers reports
  - schedules jobs
  - If the scheduler is restarted, jobs restart from where they stopped.

- **Repository** - The repository manages the persistence of Qlik NPrinting entities by saving the entities to a database.

- **Web Engine** - The web engine manages user authentication and authorization based on user roles.

- **Messaging Service** - Qlik NPrinting Messaging Service manages the communication between the Qlik NPrinting Server and the Qlik NPrinting Engines.

Qlik NPrinting Engine - Engines produce the reports based on information sent to them from the scheduler. Each engine is multi-threaded to benefit from multichannel CPUs. It is possible to install many engines on different computers to distribute the workload, and to create a high-availability reporting system. **Note**: Only one engine can be installed per computer and QlikView Desktop is required on each engine server when accessing QlikView as a data source.

Data Sources

Qlik NPrinting creates and distributes reports by using report data and visualizations from Qlik Sense and QlikView applications as data sources, which in turn are sourced from any number of data sources. Visualizations can be also created directly from within Qlik NPrinting Designer using native capabilities in third party tools such as MS Office applications and Pixel-Perfect designer.

Qlik NPrinting accesses Qlik Sense and QlikView applications running on the Qlik Sense Enterprise Server or QlikView Enterprise Server, respectively. Qlik NPrinting connects to and leverages existing Qlik Sense and QlikView resources to produce the data calculations and visualizations, leaving the actual report production for the Qlik NPrinting components described above.
Alternatively, Qlik NPrinting can store QlikView applications locally on the Qlik NPrinting server and access those for report production.
Note: Additional CPU and RAM resources are required in this scenario. The benchmark statistics in this document are based on dedicated Qlik Sense and QlikView resources.

User Access and Report Delivery

Qlik NPrinting reports can be delivered to users and user groups through email, folders, the Qlik Sense Hub, and the NewsStand portal. Users who receive their reports that are "pushed" to them in this manner are typically '**passive**' users or recipients.

Qlik NPrinting reports can also be consumed actively by users who desire to generate a new report with the latest data on the spot. This is done through Qlik NPrinting On-Demand where active users can define filter selections and 'demand' the creation of a new report. Users who demand reporting in this manner are typically '**active**' users and would usually also have a QlikView or Qlik Sense user license.

Through NewsStand, active users can also subscribe to their favorite reports by defining a recurring schedule for personal report delivery. NewsStand will also allow users to retrieve historical runs of the same report for prior period comparison.

When sizing a Qlik NPrinting system, it's important to consider the necessary report production needs for reports that will be produced and sent to passive users as well as the number of active users who will be accessing the system to demand reports. Because of both the scheduled report production and distribution as well as the ad-hoc 'on-Demand' production of reports, the underlying QlikView or Qlik Sense hosts may also need to be scaled to accommodate the increased demand for data from Qlik NPrinting.

Passive Delivery

Email – This is the most common form of report delivery. It occurs via an SMTP server accessible to the Qlik NPrinting Server.

Folders – distributing to hard drive folders avoids cluttering the inbox. In this distribution, Qlik NPrinting delivers the finished reports to hard drive folders that are accessible to the Qlik NPrinting server such as a UNC share. Folders can also be user specific. The folders serve as a staging area for recipients or other downstream consumers of the reports.

Qlik Sense Hub – When reports are sent to the hub, a recipient who has logged into Qlik Sense will see the reports alongside the streams and Qlik applications that the recipient has access to.

NewsStand - Like the hub, recipients who have logged into NewsStand will see the reports that have been sent to them. In addition, they can download previously run versions of the same report.

Active Consumption

Qlik NPrinting On Demand – When a report is enabled for "On-Demand," it can be generated outside of a pre-determined schedule by the user of a Qlik application or by any process that invokes the Qlik NPrinting OnDemand API. A Qlik user can generate the report on the fly with the option to filter the report with the current field selections in their Qlik session. An on-demand request can also be made through the RESTful API which provides a process that invokes the API with the ability to select the report, format, and filters.

NewsStand – A user of NewsStand can also subscribe to a report that the user isn't currently receiving as well as view and download the reports that were sent. In addition, a user can add a recurring schedule to the subscription defining how often the report should be refreshed and sent to the user's NewsStand inbox.

Development and Administration

Designer for developers – Qlik NPrinting Designer is Qlik NPrinting's report writing developer fat client. It must be installed on the developer's workstation and can only be invoked through the web console. The designer snaps into MS Word, Excel and PowerPoint which gives developers the ability to author reports in office formats. The designer also provides an option to author HTML as well as an advanced proprietary report editor called Pixel Perfect which offers its own charting and page design capabilities. Neither the HTML editor nor the pixel perfect editor require MS Office.

Note: Qlik Entity reports which transform a Qlik application sheet 'as is' into static document format do not require the use of the Designer.

Web Console – The web console is an administrative portal for administrators and developers. The console is where the server is configured and where reports, data connections, tasks and other content is authored and edited. Since Qlik NPrinting supports creation of custom user roles, the secure delegation of capabilities between content generation and system configuration can be defined by each administrator of Qlik NPrinting

Reports

Qlik NPrinting provides high quality reports in widely used portable formats. Qlik NPrinting Designers will design these formats using MS Office, HTML, or Qlik Pixel Perfect editor using the content and data sourced from applications hosted by QlikView or Qlik Sense. With MS office, virtually all existing native office design capabilities can be leveraged to design data from Qlik including the use of the office chart wizards. Qlik NPrinting supports common and essential reporting writing concepts such as multipage design, pagination, nested repeating elements (i.e., levels), bursting (i.e., cycling), and integration of visual, tabular and rich text content. To take advantage of all the Designer capabilities, a Qlik NPrinting designer should have both the Qlik NPrinting Designer software installed as well as a licensed local copy of MS Office.

Table: Qlik NPrinting Template Editors vs. support formats

		Qlik NPrinting Template Editors					
		Excel	**PowerPoint**	**Word**	**PixelPerfect**	**HTML**	**Qlik Entity**
Output Formats Supported	**XLS / XLSX**	✓			✓		
	PPTX		✓				
	DOC / DOCX			✓			
	PDF	✓	✓	✓	✓		
	TIFF	✓			✓		
	JPG				✓		✓
	PNG				✓		✓
	GIF				✓		✓
	BMP						✓
	HTML	✓		✓	✓	✓	
	HTM					✓	

Table: Qlik NPrinting Template Editors vs. supported reporting capabilities

		Qlik NPrinting Template Editors					
		Excel	**PowerPoint**	**Word**	**PixelPerfect**	**HTML**	**Qlik Entity**
Design Capabilities	**Images** (Qlik app Objects rendered as images)	✓	✓	✓	✓	✓	✓
	Tables (Qlik app Object dimensions and measures)	✓	✓	✓	✓	✓	
	Office design	✓	✓	✓			
	Qlik design				✓		
	Cells	✓	✓	✓	✓	✓	
	Variables (Qlik app variables)	✓	✓	✓	✓		
	Formulas	✓	✓	✓	✓		
	Pages / Pagination	✓	✓	✓*	✓*		
	Levels	✓		✓	✓	✓	

***Pagination is available through native Office capability.**

Scaling

This section focuses on the overal scaling charactersitics of the number of users, the number of cores in a server, and the number of engines within the Qlik NPrinting environment. This ensures predictable performance and scaling as an Qlik NPrinting deployment grows. As the needs increase to produce a larger number of reports, scalability becomes increasingly important.

Scaling with Users

Qlik NPrinting scales linearly with users. As the number of reports users receive grow, their response times and impact on server resources grows predictably and linearly. Of course, report creation time varies with report complexity and output type, but for a given report, its production scales linearly with users.

Qlik NPrinting reports can be delivered to thousands of users and is predictable in what impact an addition of users will have on resource consumption and report production time. The following charts show the publishing times in seconds of both QlikView and Qlik Sense reports. As defined below, the Simple report is being published for up to 5000 users and the Complex report up to 1000 users. The publishing times are shown on the primay axis on the left side and the average CPU percentage and RAM in GB usage are shown using the secondary axis on the right side.

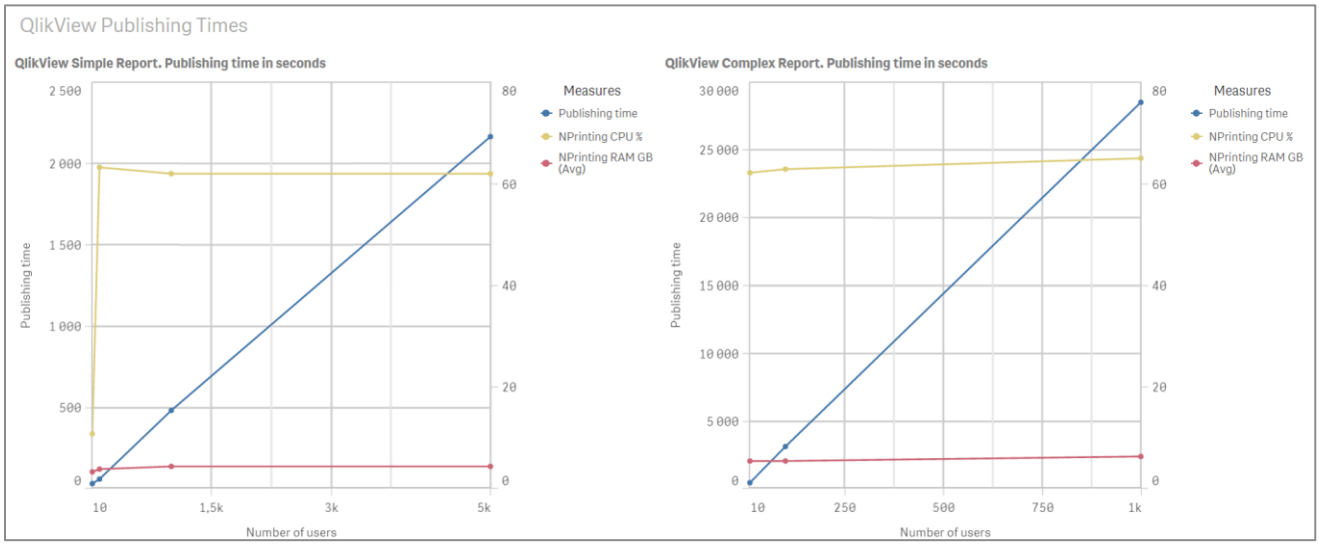

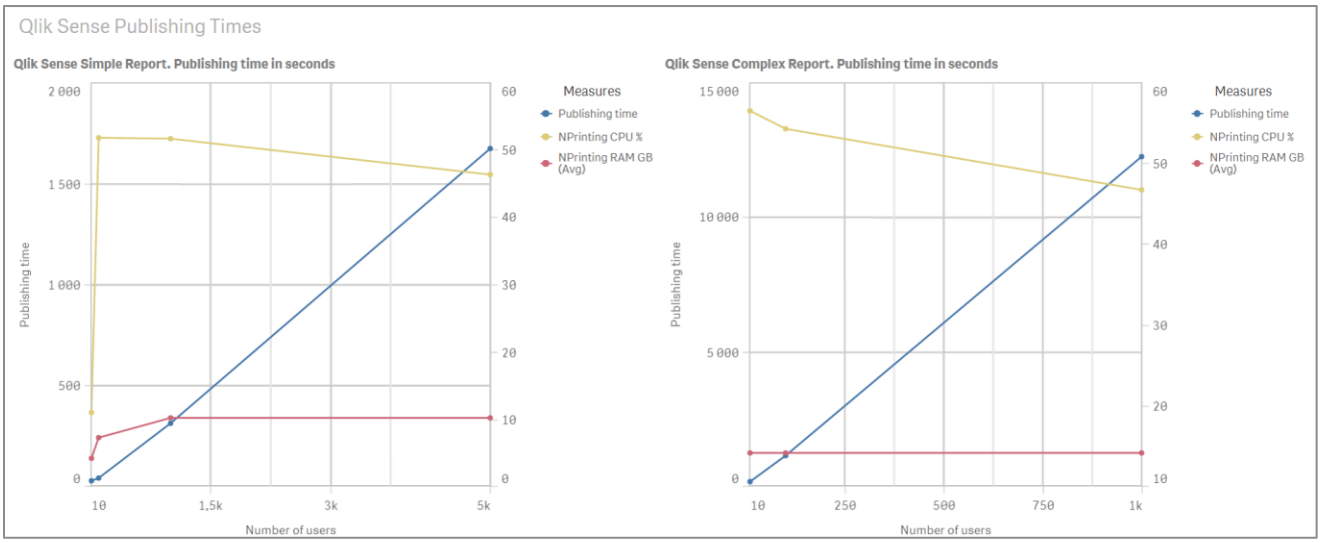

Scaling with Cores

The Qlik NPrinting reporting engine is a 64 bit, multi-threaded process that is optimized to take advantage of multiple processor cores when performing calculations. The effect that the number of CPU cores has on performance depends on whether the reports are based on connections to QlikView or Qlik Sense.

When reports are created from Qlik Sense connections, performance increases with the number of CPU cores added. The total amount of time necessary to deliver a set of reports will also depend on external factors, such as the performance of the SMTP server for example.

Note: There is a known upper limit when creating reports from QlikView connections. Performance increases with the number of CPU cores added until twelve cores are reached after which the performance increase becomes less than linear.

The following charts show test results using one XLSX medium report being published to 100 users on 4, 8 and 16 cores. In this scenario, every user has a unique filter to ensure unique report creation per user.

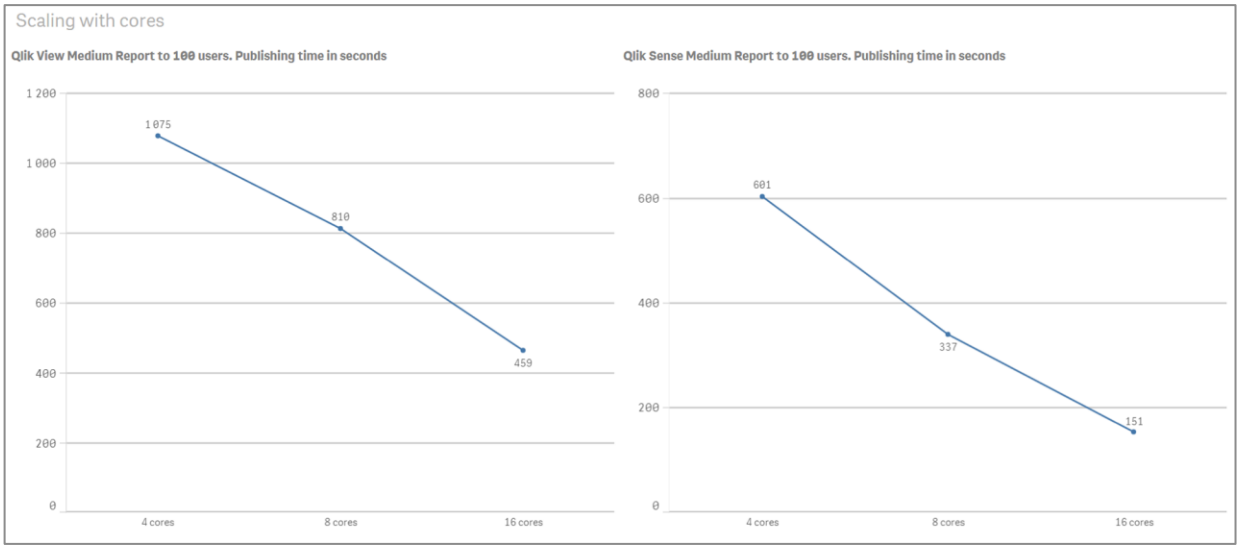

Scaling with Engines

NPrinting scales with engines. Increasing the number of Qlik NPrinting Engines associated with a Qlik NPrinting Server will result in a near linear increase in performance. For example, the time to create and deliver the same set of reports in an installation with three Qlik NPrinting engines will be almost twice as fast as an installation with one Qlik NPrinting Engine. Note: The total amount of time necessary to deliver a set of reports could also depend on external factors, such as the performance of the SMTP server for example. The following charts show the publishing time in seconds of a medium report from both QlikView and Qlik Sense to 100 users using up to three Qlik NPrinting engines

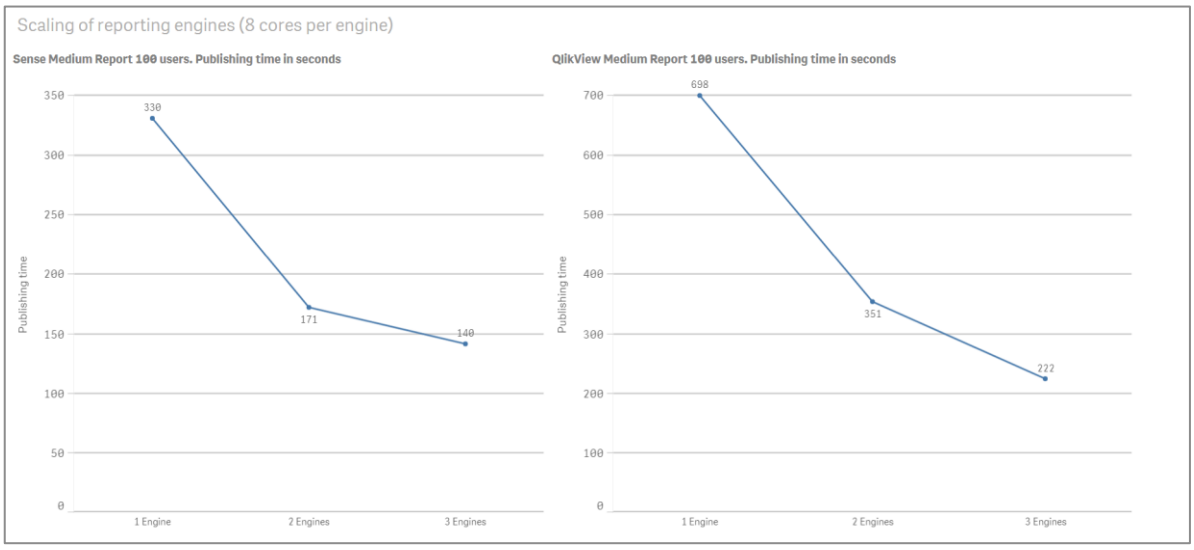

About RAM and using local QlikView documents

The amount of RAM needed on Qlik NPrinting Engine servers depends on whether QlikView documents are accessed remotely or locally. For each core of a Qlik NPrinting Engine, an instance of QV.exe is launched which opens a single QlikView document. For example, if a Qlik NPrinting Engine server has four CPU cores it means that it runs four QV.exe instances which in turn open four QlikView documents. If those QlikView documents are opened remotely on the QlikView server, then Qlik NPrinting in effect acts like a client and needs minimal RAM. (See the deployment recommendations and benchmarks below which are based on remotely accessing QlikView.) Local access of QlikView documents directly on the Qlik NPrinting server requires additional RAM to open the QlikView documents, and the document can potentially be opened one time per core.

With Qlik Sense, all connections are remote to the Qlik Sense server, so Qlik NPrinting, in effect, acts like a Qlik Sense client.

Benchmark Results

Test Setup

The tests in this paper were performed on the following software and hardware:

Service	Version	CPU	RAM
Qlik NPrinting	17.2.3 RC1	E7-4850 – 8 cores (Virtual)	32 GB
QlikView	12.10 SR1	E5-2670 – 16 cores	384 GB
Qlik Sense	3.1 SR3	E5-2670 – 16 cores	384 GB

Qlik NPrinting Reports

The test reports generated from Qlik NPrinting and used in the benchmarks are catagorized as Simple, Medium and Complex, with each produced from both Qlik View and Qlik Sense applications.

Qlik View Reports

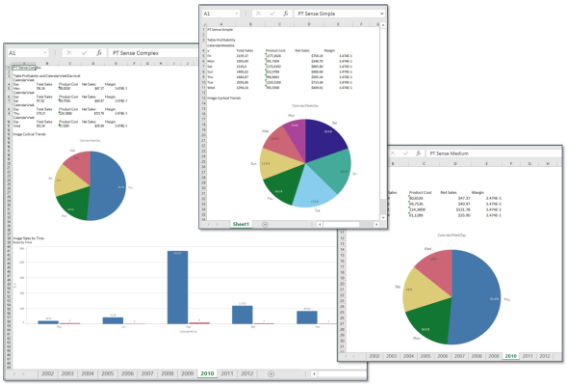

Qlik Sense Reports

Report contents:

- Simple. 1 page with 1 table and 1 image
- Medium. 10 pages with 1 table and 1 image on each page
- Complex. 10 pages with 1 table in a level and 2 images on each page

Qlik Sense and QlikView Applications

The applications that supported the Qlik NPrinting report production of the Simple, Medium and Complex reports are called PT10M, there are versions for both Qlik Sense and QlikView, with both Qlik apps each having 10 million rows of the same data. The Qlik apps are comprised of typical sales data and have several sheets showing trends such as sales and profitability across customers, products and regions in aggregate via different graphical objects such as line, bar and pie charts as well as tables.

QlikView PT10M app

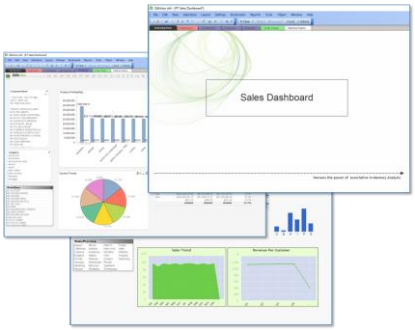

Qlik Sense PT10M app

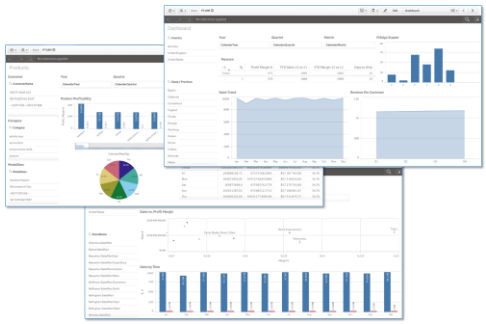

Report Throughput Results

The report design itself affects the overal report production time. In particular, the number of images, objects and levels can yield the greatest impact on the resources required to produce a report. This benchmark is based on sample reports designed to reflect common requirements that we see in our customers. Your particular reports will yield different results, but these benchmarks can provided a basis for understanding scalability and performance of Qlik NPrinting and how your deployment might relate.

Overall Performance Results

The following results are based on a 4, 8, 12 and 16 cores and 32GB RAM Qlik NPrinting server.

Report throughput results

Testcase	# of reports	Average time (sec)	seconds per report	NPrinting CPU %	NPrinting RAM GB (Avg)
Sense Simple report 17.2.3	1000	307,0	0,3	51,5	10,0
Sense Simple report 17.2.3	5000	1667,0	0,3	46,2	10,0
Sense Complex report 17.2.3	1000	12195,0	12,2	46,6	14,0
QlikView Simple report 17.2.3	1000	472,0	0,5	61,8	4,0
QlikView Simple report 17.2.3	5000	2157,0	0,4	61,7	4,0
QlikView Complex report 17.2.3	1000	28408,0	28,4	64,7	6,0

Note: The overall load on the system in terms of report generation is the same regardless of whether the report is created on-demand or via scheduling, although the on-demand process will receive a higher priority in the queue.

Core Throughput Results

The above reports were generated on 4, 8, 12 and 16 core servers to measure throughput in those scenarios, as well. The tests are summarized and averaged below to show throughput metric of reports per core per minute. The results clearly indicate that the complexity of the report is a large factor in overall throughput.

Average Core Throughput Results

Product	ReportType	Average Reports / Core / Min
QlikView	Complex	0.33
Qlik Sense	Complex	0.65
QlikView	Medium	1.26
Qlik Sense	Medium	2.93
QlikView	Simple	18.43
Qlik Sense	Simple	28.96

Deployment Examples

Below are common deployment models and sizes. Capacity ranges are calculated based on throughput results detailed in the prior section. Actual throughput will vary.

Deployment	# of Servers	Components	# of Cores	RAM
Small	1	All	4	32 GB
Medium	1	All	8	32 GB
Large	4	1 – NP Server 3 – NP Engines	12 each	32 GB each

Conclusion

Qlik NPrinting provides a secure and scalable multithreaded 64-bit reporting engine that can scale across multiple servers to produce thousands of reports for thousands of users. Qlik NPrinting allows Qlik customers to extend their investment in interactive dashboards and reach a wider audience by publishing their important business data into widely adopted portable formats such as PDF, PowerPoint, Word, Excel and others. Producing data into both Qlik dashboard as well as static formats allows companies to manage their data through a single data flow for trust and reliability across an organization.

Appendix

Complete test details for Reports per min.

NPrinting reports per min test results (full) - Average CPU percentage, Average RAM in GB

Testcase	# of Cores	Reports per min	Average time (sec)	NPrinting CPU %	NPrinting RAM GB (Avg)	QlikView Server CPU %	QlikView Server RAM	Qlik Sense Server CPU %	Qlik Sense Server RAM
Sense PT10 Simple 100 users	4 cores	128	47	32	3,5	-	-	0,5	15
Sense PT10 Simple 100 users	8 cores	261	23	25	8	-	-	0,6	13
Sense PT10 Simple 100 users	12 cores	316	19	35	10	-	-	1,1	13
Sense PT10 Simple 100 users	16 cores	400	15	28	11	-	-	0,9	14
Sense PT10 Medium 100 users	4 cores	11	552	30	4	-	-	1,8	17
Sense PT10 Medium 100 users	8 cores	22	267	38	10	-		2,8	5
Sense PT10 Medium 100 users	12 cores	40	150	36	15	-	-	2,7	6
Sense PT10 Medium 100 users	16 cores	45	132	36	17	-	-	3,3	6
Sense PT10 Complex 100 users	4 cores	4	1524	34	5	-	-	1,4	15
Sense PT10 Complex 100 users	8 cores	7	819	34	11	-	-	5,1	15
Sense PT10 Complex 100 users	12 cores	9	700	31	19	-	-	8,4	15
Sense PT10 Complex 100 users	16 cores	10	584	27	25	-	-	4,5	15
QlikView PT10 Simple 100 users	4 cores	78	77	52	2,5	0,8	16	-	-
QlikView PT10 Simple 100 users	8 cores	188	32	57	2,5	1,7	14	-	-
QlikView PT10 Simple 100 users	12 cores	214	28	44	3	2	14	-	-
QlikView PT10 Simple 100 users	16 cores	207	29	24	3	1,4	14	-	-
QlikView PT10 Medium 100 users	4 cores	7	916	73	3	1,2	20	-	-
QlikView PT10 Medium 100 users	8 cores	11	557	64	3	1,9	20	-	-
QlikView PT10 Medium 100 users	12 cores	15	410	51	3	2,5	21	-	-
QlikView PT10 Medium 100 users	16 cores	13	459	35	4	2,3	21	-	-
QlikView PT10 Complex 100 users	4 cores	2	2991	76	2,5	1,2	17	-	-
QlikView PT10 Complex 100 users	8 cores	3	1906	61	3	3,8	17	-	-
QlikView PT10 Complex 100 users	12 cores	3	1881	42	4	1,8	18	-	-
QlikView PT10 Complex 100 users	16 cores	3	2115	30	4	1,6	15	-	-