

WHITE PAPER

Qlik NPrinting Scalability

TABLE OF CONTENTS

| | |
|---|-----------|
| Summary | 2 |
| About Qlik NPrinting | 3 |
| Architecture | 4 |
| QlikView® vs Qlik Sense® Sources | 4 |
| Test Setup | 5 |
| Test Hardware | 5 |
| What is a Core? | 5 |
| Test Reports | 6 |
| Test Qlik Sense and QlikView Applications | 7 |
| Test Users | 7 |
| Results | 7 |
| Cores | 7 |
| Recipients | 10 |
| RAM | 10 |
| Additional Tests | 11 |
| NPrinting® Engines and Qlik Sense Servers | 11 |
| Report Quality | 12 |
| Cycling | 12 |
| Storage | 13 |
| Deployment Examples | 14 |
| Conclusions | 15 |

SUMMARY

Qlik NPrinting is the reporting platform for QlikView and Qlik Sense. It lets you create reports from your QlikView documents and Qlik Sense apps and distribute them automatically in a range of standard formats. Among the most common questions we get from enterprises is how does NPrinting scale? Highlighted below, and in the details that follow, are what you need to know about NPrinting multi-engine enterprise scalability:

- Our Qlik NPrinting report publishing scales in a predictable fashion.
- **NPrinting Engines & Engine Cores:** When you add more logical cores and more NPrinting Engines, you improve report publishing times and report publishing capacity.
- **Maximum Engine Cores with QlikView Sources:** Increasing the number of cores on the NPrinting Engines increases the report publishing capacity of an NPrinting site to a maximum of 12 logical cores per NPrinting Engine when using QlikView applications as a reporting source.
- **Maximum Engine Cores with Qlik Sense Sources:** Increasing the number of cores on the NPrinting Engines increases the report publishing capacity of an NPrinting site to a maximum of 16 logical cores per NPrinting Engine when using Qlik Sense applications as a reporting source.
- **NPrinting Engine RAM:** More RAM is required on the NPrinting Engines when deploying NPrinting Engines with more cores.
- **Report Complexity:** Higher report complexity increases the publishing time of the NPrinting report.
- **Report Quality:** Higher Report Quality settings increase the publishing time of NPrinting reports that are based on Qlik Sense sources.
- **QlikView and Qlik Sense Server Resources:** Deploying larger NPrinting Engine servers to increase report publishing capacity depends on ample availability of hardware resources on the QlikView and/or Qlik Sense server(s) that are used as sources for NPrinting reports.
- Avoid using QlikView local connections on production NPrinting sites.
- Report publishing performance with our Qlik NPrinting November 2018 is generally 30-50 percent more efficient, and in some cases, up to twice as efficient as Qlik NPrinting version 17.2.3.

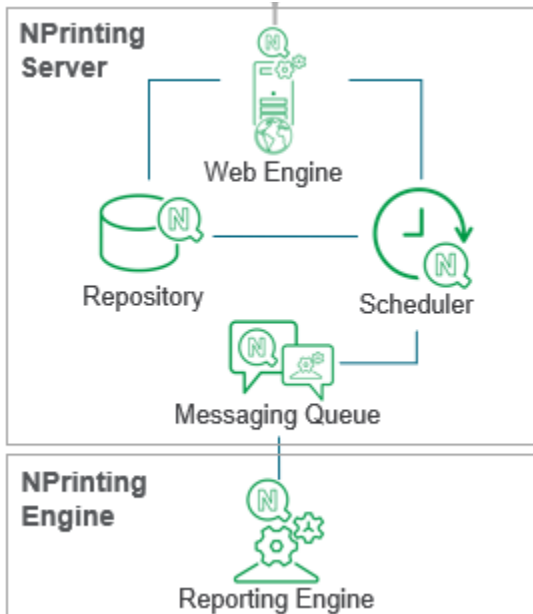
ABOUT QLIK NPRINTING

Qlik NPrinting is our scalable reporting platform that produces high-quality reports in common portable document formats.

Qlik NPrinting uses QlikView and Qlik Sense applications as the sole data sources for generating reports. This ensures efficient and trusted reuse of Qlik applications for your reporting needs. Since version 17, we have delivered Qlik NPrinting as a multi-threaded solution that can be deployed over multiple nodes to meet your enterprise report production solution needs. You can use Qlik NPrinting to distribute reports to recipients through email, folders, and online via Qlik NPrinting NewsStand or the Qlik Sense hub. Qlik NPrinting also supports dynamic interactive reporting with Qlik NPrinting On-Demand, allowing your users to adjust data filters and produce reports whenever they need them. All report recipients and your users must be preregistered in Qlik NPrinting or synced from a third-party user repository such as LDAP.

Architecture

Qlik NPrinting consists of the following components:



Qlik NPrinting Web Engine

- Web entry point into NPrinting
- Manages authentication and authorization

Qlik NPrinting Repository

- Centralized persistence of Qlik NPrinting content

Qlik NPrinting Scheduler and Messaging Queue

- Task scheduling, prioritization, distribution, and delivery

Qlik NPrinting Reporting Engine

- Generates reports from QlikView / Qlik Sense
- Can be separated from the NPrinting server
- Multiple Engines supported

A Qlik NPrinting site includes an NPrinting server and one or more NPrinting Engines. Our NPrinting server plays a largely administrative function, managing web logins to the NPrinting Web Console, NPrinting NewsStand, and NPrinting On-Demand, as well as managing your report queue, report scheduling, and storage.

NPrinting Engines produce reports and are your main scale point for increasing report publishing. When scaling NPrinting, the number of cores and number of Engines you use directly impacts the number of reports you can produce in a set time frame.

QlikView vs. Qlik Sense Sources

When you generate reports from QlikView, NPrinting Engines use a local installation of QlikView Desktop to form local and server connections to Qlik Application documents (QVWs). NPrinting Engine will open multiple instances of QlikView Desktop, allowing it to connect multiple times to one or more of your QlikView applications so it can source data and produce your reports more efficiently. This helps NPrinting resolve content required to generate one or more concurrently executing reports.

Each thread is referred to as a **content resolver**, and by default, the maximum number of content resolvers controlled by a single NPrinting Engine is equivalent to the number of logical cores present on

the NPrinting Engine. This means an NPrinting Engine with 12 logical cores may operate up to 12 instances of QlikView Desktop – each being used to extract content and data from one or more of your QlikView applications to produce your report. When using QlikView local connections, each instance of QlikView Desktop will open one or more applications locally in RAM, driving up hardware requirements. Due to the relatively high-resource needs of local connections, we recommend you use QlikView server connections or QlikView cluster connections for your production needs.

When you use Qlik Sense applications as sources for NPrinting, there is no need for a local desktop client because all connections are made over HTTPS. We recommend sourcing reports from Qlik Sense applications to avoid the overhead that comes with launching the local desktop client. For this reason, NPrinting allows double the number of content resolvers to be operated by a single NPrinting Engine when sourcing from Qlik Sense applications. For example, an NPrinting Engine with 12 logical cores may operate up to 24 concurrent HTTPS connections to Qlik Sense – each being used to extract content and data from one or more Qlik Sense applications for report production.

QlikView vs Qlik Sense

Connectivity to Qlik Sense applications over HTTPS requires less overhead than QVP connections to QlikView applications.

By default, NPrinting is configured to allow twice the concurrent connection threads (or content resolvers) to Qlik Sense applications than QlikView.

You can generate NPrinting reports from both QlikView and Qlik Sense in the same NPrinting site. Since NPrinting will dedicate distinct content resolvers for both QlikView and Qlik Sense, we recommend you deploy NPrinting servers with more RAM and faster clock speed when deploying NPrinting servers that will be used to generate reports from both QlikView and Qlik Sense at the same time.

Test Setup

In 2019, Qlik had automated benchmark tests conducted in a multi-server lab test environment, these were the main variables tested:

- # of recipients
- Report complexity (simple / medium / large)
- # NPrinting Engine cores
- QlikView vs Qlik Sense sources

Additional tests examined the impact of:

- Report quality / # NPrinting Engines / Cycling

Test Hardware

Virtualized servers were used for our benchmark tests. During testing, the number of cores was increased from 4 to 16 to understand the impact on report publishing performance.

Note: Additional boundary testing (up to 36 cores) was performed separately.

| Product | Version | Server | RAM |
|------------|---------------|---------|-----|
| NPrinting | November 2018 | E7-4850 | 32 |
| QlikView | November 2017 | E5-2670 | 384 |
| Qlik Sense | June 2018 | E5-2670 | 384 |

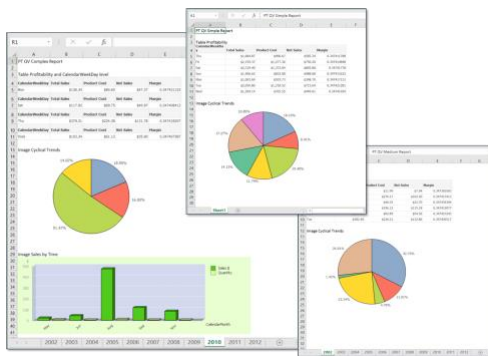
What is a Core?

In this paper, a core is deemed to be a *logical* core as defined by the configuration of the operating system. A Windows instance with 8 physical cores – without hyperthreading enabled – will provide 8 logical cores to NPrinting. The same instance – with hyperthreading enabled – will provide 16 logical cores to NPrinting. The benchmark results and recommendations provided here pertain to logical cores.

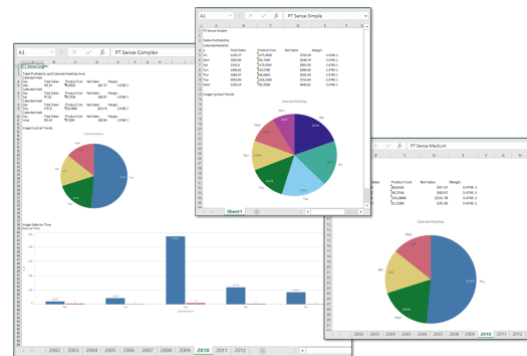
Test Reports

Test reports generated from NPrinting and used in our benchmarks were categorized as simple, medium, and complex. Each was produced from both QlikView and Qlik Sense applications.

QlikView Test Reports



Qlik Sense Test Reports

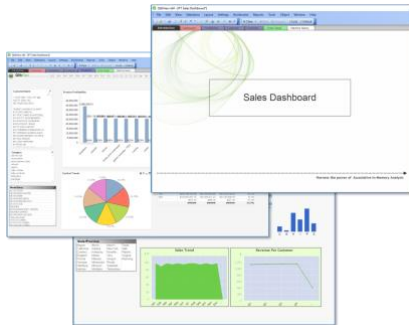


- Simple: 1 table and 1 image (Single-sheet Excel report)
- Medium: 1 page element with 1 table and 1 image on each page (11-sheet Excel report)
- Complex: 1 page element with 1 table in a level and 2 images on each page (11-sheet Excel report)

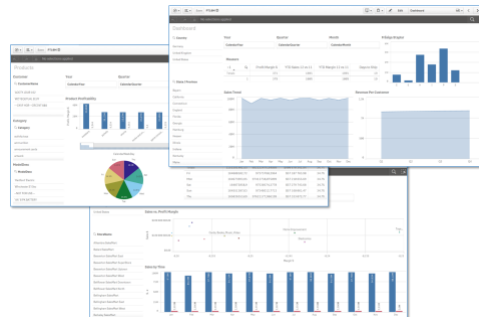
Test QlikView and Qlik Sense Applications

The QlikView and Qlik Sense apps used as sources for our testing were comprised of typical sales data. They had several sheets showing trends – including 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. Both apps contained 10 million rows of the same data.

QlikView Test app



Qlik Sense Test app



Test Users

Testers loaded recipients into NPrinting through batch scripts. Primary tests were done in buckets of 100 and 500 recipients.

Note: Additional boundary testing (up to 10,000 recipients) was performed separately.

Results


Our benchmark testing revealed insights that can help you optimize the use of NPrinting cores, number of recipients, and amount of RAM.

Cores


Reports per minute increased linearly with the addition of more NPrinting Engine cores. This was the case for simple, medium, and complex reports, as well as when generating concurrent reports of both QlikView and Qlik Sense.

The following results were achieved when producing a single report for 100 recipients, where each recipient had a different filter:

Qlik Sense – NPrinting Engine Cores – Reports Per Minute




QlikView – NPrinting Engine Cores – Reports Per Minute




Increasing NPrinting Engine cores showed clear gains in publishing capacity with some reduction in efficiency as cores were increased. This is depicted graphically in **reports per core per minute**.

Qlik Sense – NPrinting Engine Cores – Reports Per Core Per Minute


Sense PT10 Simple reports/min/core



Sense PT10 Medium reports per min




Sense PT10 Complex reports/min/core




QlikView – NPrinting Engine Cores – Reports Per Core Per Minute


QV PT10 Simple reports/min/core



QV PT10 Medium reports/min/core



QV PT10 Complex reports/min/core



Not shown, but notable, was how productivity decreased more rapidly above 12 cores with QlikView.

With Qlik Sense, diminishing returns to report publishing occurred after 16 cores. See the Complex Reports using NPrinting April 2018 Publishing Time graph. After 28 cores, there was no improvement to publishing time.

We recommend NPrinting Engines be deployed to a maximum of 12 cores when using QlikView as a source and to 16 cores when using Qlik Sense as a source.


Adding Cores to NPrinting Engine(s)

Adding cores increases the potential reporting capacity of NPrinting.

There are diminishing returns to report publishing capacity above 12 cores when using QlikView applications and above 16 cores when using Qlik Sense applications as sources.

Qlik Sense – NPrinting Engine Cores – Publishing Time

Sense Complex Report 1000 and 500 recipients. Publishing time in seconds



Recipients



Our setup tested 100 recipients versus 500 recipients. We maintained **reports per minute** when comparing 500 recipients versus 100 recipients on NPrinting Engines with both 4 and 16 cores, across simple, medium, and complex reports with both QlikView and Qlik Sense.

Adding Recipients


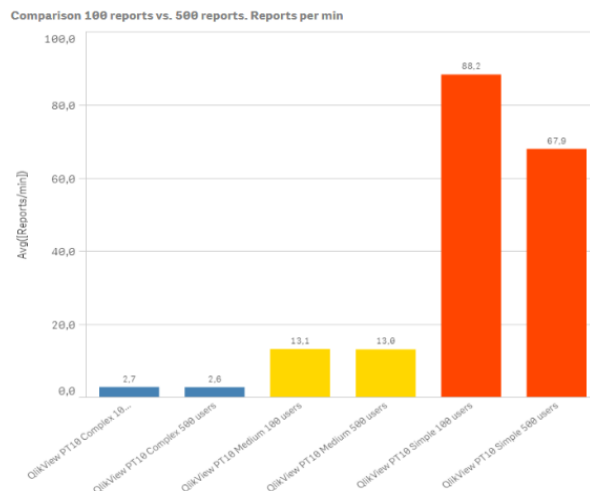
NPrinting maintains report production with higher numbers of recipients.

NPrinting has been tested and can produce reports for up to 10,000 recipients.

Qlik Sense – Recipients – Reports Per Minute



QlikView – Recipients – Reports Per Minute



RAM

When more cores are present on an NPrinting Engine, the report publishing capacity of the Engine increased, driving up the need for more RAM. We recommend deploying NPrinting Engines with 12 or more cores with a minimum of 64GB of RAM.

Additional Tests

Additional tests were run to determine the impact of NPrinting Engines, additional Qlik Sense servers, report quality, and cycling to a high number of recipients.

NPrinting Engines and Qlik Sense Servers

Reports per minute increased when an additional 12 core / 32GB RAM NPrinting Engines were added. After 3 engines and with a larger number of recipients, the Qlik Sense source server became saturated, prolonging report publishing.

When rerunning the test with a dual-node Qlik Sense site, the saturation of Qlik Sense only occurred on 5–6 Engines for more complex reports and a higher number of recipients.


Note: These Engine and Qlik Sense node benchmark tests were performed using NPrinting April 2018 and Qlik Sense February 2018.

When scaling complex reports to a higher number of cores, Engines, and recipients, we recommend increasing the capacity of the underlying Qlik Sense source servers to allow for increases in report publishing capacity.


Qlik Server Saturation

At higher levels of report publishing, Qlik Sense and QlikView server hardware can be a bottleneck and must be scaled to meet the needs of NPrinting. QlikView and Qlik Sense servers can be scaled by additional RAM and CPUs or by adding more servers.

Qlik Sense – Single Node Site – Reports Per Minute



Qlik Sense – Dual Node Site – Reports Per Minute




Report Quality

You can set the quality of NPrinting reports higher in publish tasks, as well as for on demand. High or Very High report quality settings improve the resolution in published NPrinting reports. Higher-quality reports can take longer to produce. It is also worth noting that higher-quality reports often result in larger document sizes on disk.

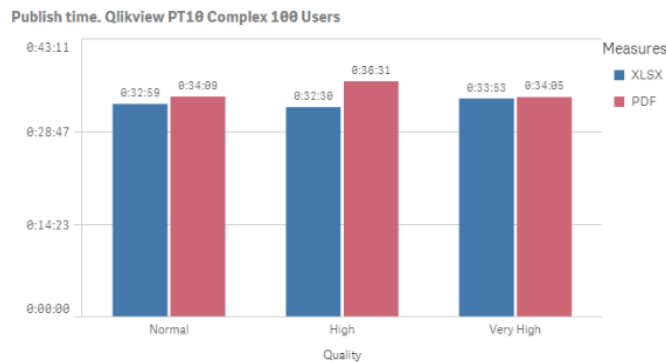
In our testing, the publish time to generate reports increased with higher-quality reports vs lower-quality reports when using a Qlik Sense app as the source. With QlikView, the publish time did not increase.

Note: These tests were performed with NPrinting June 2018, Qlik Sense February 2018, and QlikView November 2017.

Qlik Sense – Report Quality – Publish Time



QlikView – Report Quality – Publish Time



Cycling


Cycling splits report execution across one or more dimensions, producing discrete, filtered outputs for all possible filter combinations found in the dimension(s). For example, year, quarter, product line, country, etc.

Cycling is commonly used to batch produce a complete set of reports for all filter combinations. The cycle distribution is often done to hard-drive locations (folders) employing NPrinting as a reporting engine only. However, cycling can also be distributed to NewsStand or email. Recipients only receive copies of reports consistent with their user filters.

In testing, we increased a cycle to produce up to 10,000 discrete reports with a single NPrinting server and single Qlik Sense server. Publish times increased in a linear fashion up to 10,000 recipients. With complex reports, additional hardware was required to reach 10,000 recipients.

NOTE: This test was conducted with NPrinting September 2018 and Qlik Sense June 2018.

Qlik Sense – Cycles – Publish Time



Storage

NPrinting stores application metadata, report templates, and queue information on the NPrinting Server. How much storage is required depends on the # of recipients, tasks, and reports that have been specified, but is not typically greater than 10GB.

You will need additional storage for storing your reports, depending on the mode of distribution and length of time. NewsStand and folders can provide access to historical reports for many months. Most reports are under 1MB in size, but size increases with the number of pages and the use of higher-quality report settings.

Storage Locations

| NPrinting Distribution | Storage Location |
|------------------------|--------------------------------------|
| Email | <not stored> |
| NewsStand | NPrinting Server |
| Folder | Folder Location of Choice |
| On-Demand | NPrinting Server (1 month) |
| Qlik Sense Hub | Qlik Sense Shared Persistence Folder |

Deployment Examples

Deployment models and sizes vary. The most common are in the table shown. Capacity ranges are calculated based on throughput results (detailed in the prior section). Actual throughput will vary.

| Deployment Size | # of Servers | # of Cores | RAM | Storage* |
|-----------------|--|------------|-----------|----------|
| Small | 1 NPrinting Server with Engine | 4 | 32GB | 25GB+ |
| Medium | 1 NPrinting Server with Engine | 8 | 32GB | 50GB+ |
| Large** | 1 NPrinting Server with 3 separate Engines | 12 each | 64GB each | 100GB+ |

*You will need more storage when maintaining historical reports in NewsStand or on shared folders, or when using NPrinting On-Demand. NPrinting On-Demand stores reports for a month.

**With large deployments, you will need to ensure adequate resources on the source QlikView and/or Qlik Sense servers, especially for a high # of recipients, high # of reports, large cycles, and very high report quality settings.

Conclusion

Our NPrinting solution is an extension of our Qlik platform, allowing you to reuse trusted Qlik applications for reporting needs. Recent improvements have made NPrinting report publishing up to twice as efficient, and in general, 30–50 percent more efficient than NPrinting 17.2.3.

With additional hardware, Qlik NPrinting will publish more while maintaining an efficient rate of publishing to small and large groups of recipients. Report complexity, report quality settings, and the use of QlikView versus Qlik Sense as source applications additionally factor into the rate of report publishing.

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



© 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. Other marks and logos mentioned herein are trademarks or registered trademarks of their respective owners. NPRINTINGSCALEWP032619_DB