

Kentico 8

Performance

June 2014

Table of Contents

Disclaimer.....	3
Executive Summary.....	4
Basic Performance and the Impact of Caching.....	4
Database Server Separation	5
Web Farm Usage	5
Impact of Site Size on Performance	6
Cloud Computing.....	6
How Our Tests Were Conducted.....	8
Hardware Configuration.....	8
Software Configuration	9
Kentico 8 Settings.....	9
Load Test Settings.....	10
Performance Test Results	11
Pages Views per Hour.....	11
Average Page Time	12
Processor Time	13
When Output Caching Doesn't Help	14
Further Reading.....	14

Disclaimer

This report was conducted by Kentico Software with the intention of informing customers about performance in Kentico 8. Kentico has done its best to ensure an unbiased test. However website performance depends on many things, such as computer hardware, network configuration, client configuration, operating system and software configuration, site content, number of items in the Kentico database, information architecture, custom code, and other factors. Kentico Software cannot, therefore, guarantee that the same results will be achieved with any other than the tested configurations. The reader of this report uses all information contain within at their own risk. Kentico Software shall in no case be liable for any loss resulting from the use of this report.

Executive Summary

Kentico 8 for ASP.NET provides excellent performance and scalability. Being built on the Microsoft ASP.NET platform, it leverages all its power.

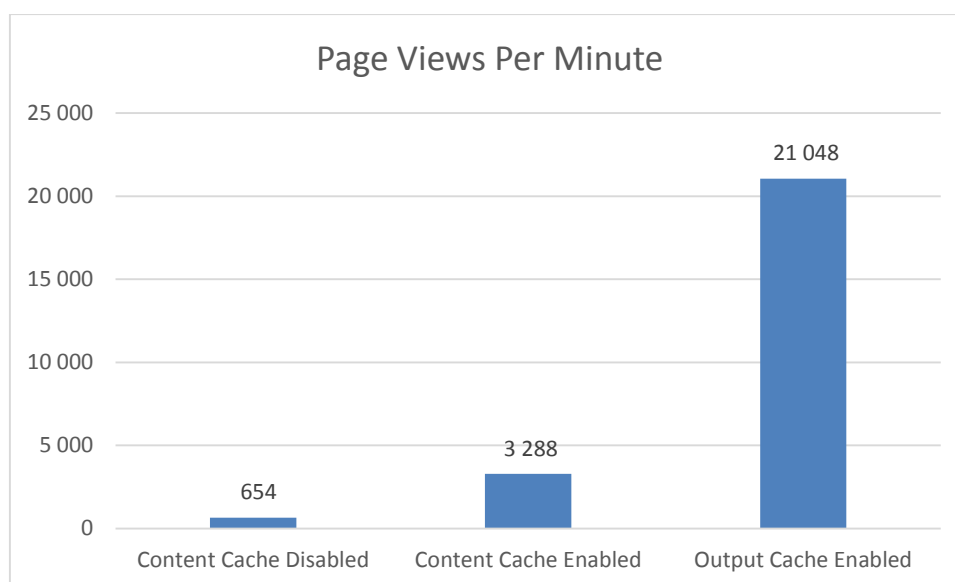
Tests were performed internally by Kentico staff on very common hardware (Intel Core 2 Quad at 2.66 GHz, 10k SATA II disks, 8 GB memory). Kentico has not used any high performance servers, so results may be even better on more powerful hardware.

Basic Performance and the Impact of Caching

The slowest parts of a web application are typically accessing the database and rendering the content for a web browser. Kentico 8 optimizes performance by storing content that is frequently accessed in a server memory. This mechanism is called **caching**. When another visitor comes to the same page, the page is already stored in the very fast computer memory and Kentico 8 can quickly send it to the browser without repeatedly accessing the database and rendering the page.

Caching can be configured for a particular part of the page – this is called **content caching**. You can also configure **output caching** which stores the whole page pre-rendered in the memory.

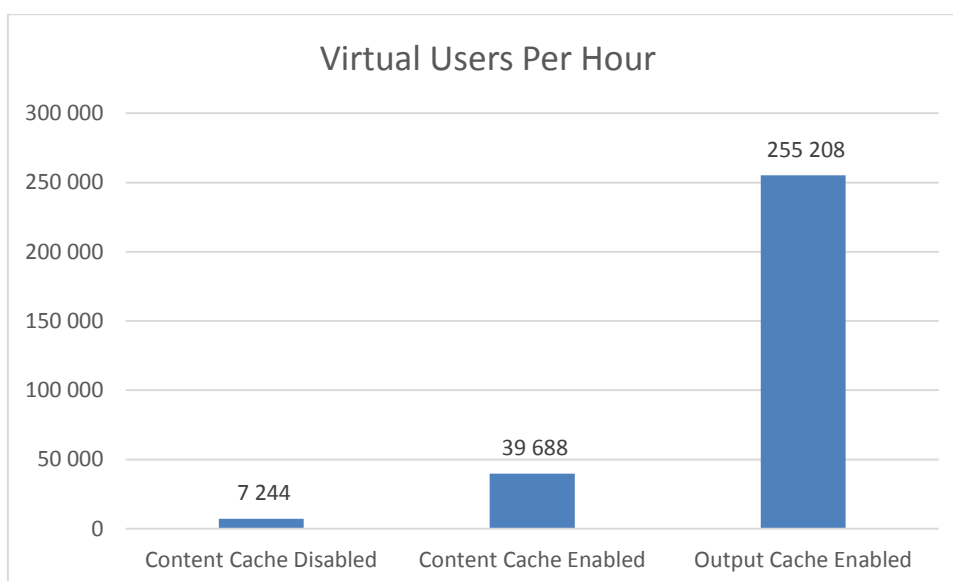
The following graph shows the impact of caching on overall performance. The values represent the number of pages viewed by virtual users per minute.



A total of 654 page views per minute with content cache disabled means about 40 thousand page views per hour, which means almost 1 million page views per day. That's about **30 million page views per month**. If you enable the content cache, you can handle 5

times more traffic, i.e. almost **150 million page views per month**. In this scenario, with the output cache enabled, Kentico 8 can handle almost **1 billion page views per month**.

Of course, this depends heavily on the content of your website, the features you use, and the hardware you have available. To improve performance, you can use CDN (Content Delivery Network) services to serve static files like images and relieve the web server of client requests. You should also use client caching which enables the user's browser to cache static resources so it doesn't have to send requests for these files to the web server at all.




As you can see in the graph above, during our **one-hour test** with the output cache enabled, more than **250,000 unique visitors** each visited 5 pages of our Sample Corporate site News section, which meant serving more than **1.25 million page views per hour**.

Database Server Separation

If you need high performance, it's recommended that you install the web server and the database server on two different machines. This allows you to distribute the computing and achieve shorter response time, especially if you cannot use caching.

Web Farm Usage

Web farms allow you to distribute computing among **multiple web servers** that all provide the same content. This gives you **high scalability** and it can also be used for achieving **high availability** of your site – if one web server in the web farm stops working, the other server(s) will serve the content instead of the broken server.



Performance **improves proportionally with the number of web servers**, especially in scenarios where site visitors mainly read content that has been properly cached. But there are some scenarios in which adding more web servers doesn't help much, especially where there are a lot of create/update/delete operations and synchronization between web servers is needed or where caching can't be used and the database server is not powerful enough.

Impact of Site Size on Performance

Website performance also depends on the number of pages and other items in the Kentico database. Kentico 8 was optimized for a high number of items and was tested with **100,000 documents** stored in the database. Not only can public-facing websites handle this number of items without negative effect on usability or user interface responsiveness, but so can the Kentico 8 administration interface.

Of course, Kentico must be properly set up and optimized for such scenarios. As each project is unique, this could require some customization and optimization of the SQL queries and database indexes. Performance of the live site with caching enabled is practically the same regardless of site size, except for the first page load.

Cloud Computing

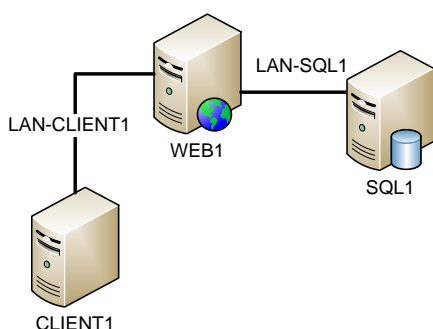
Kentico 8 can also be used with cloud computing platforms, which provide both high performance and availability. Kentico 8 natively supports Microsoft Azure (including Microsoft Azure Storage) and Amazon EC2 (including Amazon S3 Storage), and we have numerous clients who use Kentico on these platforms. Our customers also use other cloud platforms, such as Rackspace Cloud.

Kentico 8 can be run on multiple instances in both Microsoft Azure and Amazon EC2, which makes you ready for virtually unlimited traffic. You can also leverage their CDN capabilities to optimize the delivery of uploaded files, such as images, video, PDF documents, etc.

How Our Tests Were Conducted

Hardware Configuration

As you can see in the following image, we used 3 computers dedicated to performance testing. Client1 was the computer that generated the load to the web server, called Web1. Sql1 was the database server for storing databases used by Web1.



Motherboard: Giga-Byte EP45C-DS3R (Intel® Socket 775, Intel® P45)

Processor: Intel® Core™2 Quad Q9400 BOX (2.66GHz)

Memory: Kingston 2G-UDIMM (PC2-6400)

Capacity	8 GB (Kit 4x 2 GB)
Type	DDR2-SDRAM
Memory Bus	1066 MHz

HDD: 1x Western Digital WD1500HLFS VelociRaptor 5RZ

Format	3.5"
Capacity	150 GB
Interface	SATA II
Speed	10000rpm
Access Time	4.2ms / 4.7ms
Cache	16 MB

LAN: Realtek PCIe GBE Family Controller – 1 Gbit/s

These components were the same for all the servers. The only difference was that Sql1 had more memory installed (16 GB).

Software Configuration

- Microsoft Windows Server 2012 R2
- Microsoft Internet Information Services 8.5
- Microsoft SQL Server 2014
- Microsoft Visual Studio 2013 (Ultimate Edition)

Kentico 8 Settings

CONFIGURATION FOR ALL TESTS

- Enable output compression: **True**
- Use progressive caching: **True**
- Client cache (minutes): **100**
- Allow JavaScript minification: **True**
- Allow CSS minification: **True**
- Disable debugging: **True**
- Enable on-line marketing: **False**
- Enable A/B testing: **False**
- Enable multivariate testing: **False**
- Enable content personalization: **False**
- Enable web analytics: **False**

CACHING CONFIGURATIONS

- **Content Cache Disabled:** Page info (page data and metadata) cache enabled for 100 minutes, content caching set for 0 minutes
- **Content Cache Enabled:** Content cache (web part/control-level caching) enabled for 100 minutes¹
- **Output Cache Enabled:** Output cache (whole page cached in memory) enabled for 100 minutes

¹ We used default Output filter settings. The use of additional XHTML filters decreases performance by around 30% when using content caching – this can be solved by writing XHTML-compliant code so that the XHTML filter doesn't have to be used.

Load Test Settings

Tests were based on visits to the News section of our Sample Corporate site. There were more than 100,000 documents in the database to demonstrate the power of caching.

- All tests were performed for a duration of 1 hour (+ 10 seconds warm-up time)
- Domain name was in the host file (localhost was not used)
- Load Tests were executed in Visual Studio 2013 (Ultimate Edition)
- The constant load of 50 simultaneous virtual users was used (100% new users)
- When a virtual user finished visiting their 5 page requests, a new virtual user was created right away
- There was no “Think Time” between requests
- Dependent requests were also parsed
- No HTTP, DNS, Socket, Windows or Kentico 8 errors were allowed during tests
- Kentico application restarts were not allowed during tests
- IIS was restarted before every test execution and the application was warmed up with a few requests

KENTICO 8 – SAMPLE CORPORATE SITE

~/News.aspx

~/News/Apple-iPad-2-In-Stock.aspx

~/News/Community-Website-Section.aspx

~/News/Company-Growth-Exceeds-Expectations.aspx

~/News/New-Consulting-Services.aspx


PERFORMANCE COUNTERS WATCHED

- Processor(_Total)\% Processor Time
- Process(w3wp)\Private Bytes
- .NET CLR Memory(w3wp)\# Bytes in all Heaps
- .NET CLR Memory(w3wp)\# Gen 0 Collections
- .NET CLR Memory(w3wp)\# Gen 1 Collections
- .NET CLR Memory(w3wp)\# Gen 2 Collections
- .NET CLR Memory(w3wp)\% Time in GC
- .NET CLR Exceptions\# Exceps thrown / sec
- ASP.NET\Requests Queued
- Memory\Available Mbytes
- PhysicalDisk(_Total)\Disk Read Bytes/sec
- PhysicalDisk(_Total)\Disk Write Bytes/sec
- Network Interface\Bytes Received/sec
- Network Interface\Bytes Sent/sec

Performance Test Results

Pages Views per Hour


The “Page View” statistic shows how many complete pages are loaded, including dependent requests such as images, javascript files, stylesheets, etc. So, the results for Requests per Second (RPS) would be much higher (proportional to the number of dependent requests). On the other hand, some of these files could be cached on the client side and therefore wouldn’t be requested more than once per user session. In the following image, you can see the number of Page Views per hour for our three cache scenarios.



In one hour, with output cache enabled, besides the **1.2 million main requests**, there were another **8.1 million dependent requests performed** and another **11.2 million dependent requests cached**.

Average Page Time

The “Page Time” statistic includes the response time of all dependent requests. As you can see in the following graph, the loading of pages that display documents with caching disabled is very slow – it takes several seconds for a high number of documents stored in the database. In contrast, output caching only has to take the HTML code cached in the memory and send it to the browser, so it's extremely fast.

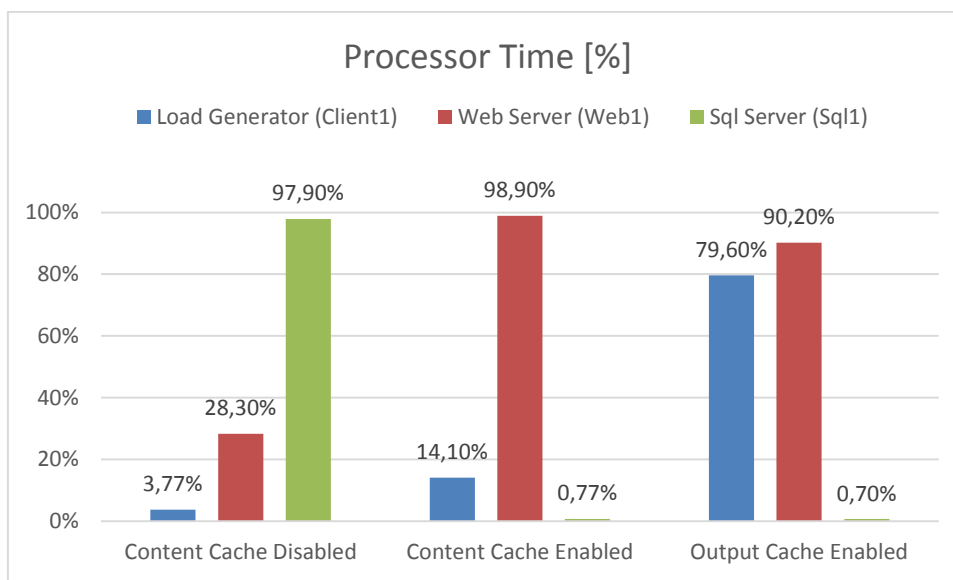


There were more than 100,000 documents in the database. The News.aspx page displayed teasers of the top five news documents and the other pages displayed a specific news document.

Page 0	~/News.aspx
Page 1	~/News/Apple-iPad-2-In-Stock.aspx
Page 2	~/News/Community-Website-Section.aspx
Page 3	~/News/Company-Growth-Exceeds-Expectations.aspx
Page 4	~/News/New-Consulting-Services.aspx

Processor Time

The “Processor Time” performance counter shows the hardware limitation of both the system and the tests. As you can see in the following graph, the processor time is a bottleneck in each scenario but on different computers.



With content cache disabled, the load generator is almost idle (doesn't have to generate any more requests), the web server has plenty of resources available, while the SQL server is overloaded. This is the scenario in which it makes perfect sense to separate the database from the web server. If it was a shared server, with both website and database, performance would be significantly lower.

With content cache enabled, the load generator has to generate more requests to the web server because the web server is able to process more of them. Notice that the SQL server is idle, since the content is cached in the application memory. The web server is now a bottleneck.

When the output cache is enabled, the load generator isn't able to generate enough load to fully test the power of the web server. If we had a more powerful load generator, the performance results of the web server would be even better.

When Output Caching Doesn't Help

As you can see from the test results, the use of the output cache may increase overall performance significantly. But it's important to say that in some cases the output cache should not be used:

- **Content that changes more often than every minute** – if you need to display real-time data, you cannot use output caching.
- **Personalized content with a high number of users** – if you need to display content personalized to the current user, such as user name, or if you need to restrict displayed content by the current user's permissions, the output cache cannot be used without some drawbacks. Kentico 8 can cache such pages, but since the page is stored in the memory for each authenticated user, it may consume lots of memory.

You can still use content caching for chosen web parts/controls in such cases, but it provides a significantly lower performance.

Further Reading

You can find more information about tweaking Kentico performance in the following materials:

- <https://docs.kentico.com/display/K8/Optimizing+website+performance>
- <https://docs.kentico.com/display/K8/Configuring+caching>
- <http://devnet.kentico.com/articles/performance-troubleshooting-tips-and-tricks>