



“Now that we've added VCS to our web-based toolbox, we can see the cache hit rate for different hosts and publications, and also drill down into hits and misses for these different sites and all their subpages”

Case Study:

Digital media: vg.no

Varnish Custom Statistics (VCS)

How Norway's most popular online news site uses Varnish VCS to monitor web traffic and resolve problems in real-time.

Background

If your website's traffic flows through Varnish Cache then this is the ideal place to gather data about your traffic.

However, gathering and making sense of all this data can be a challenge. Varnish Custom Statistics (VCS) is one of the latest additions to Varnish Software's toolkit and it is designed to help you do this. It captures and provides real-time aggregated statistics on all the traffic that goes through Varnish Cache. It allows users to define their own rules, so-called keys, in order to detect and collect data on their websites' traffic. This enables them to identify and categorise problematic patterns in real time and to react appropriately as it happens. Using VCS allows you to keep your website healthy and running with optimal performance, as you are able to track and analyse data in detail and in real time, thereby preventing any problems.

Varnish Software's very first customer¹, the Norwegian online newspaper vg.no, has been using VCS since it was first launched. After having used VCS since 2012 they now regard it as an essential tool for discovering bugs and problematic behaviour on their web publications. "We are running Varnish in front of all of our publications, six Varnish installations in total. Now that we've added VCS to our web-based toolbox, we can see the cache hit rate for different hosts and publications, and also drill down into hits and misses for these different sites and all their subpages," said Audun Ytterdal, Head of IT Operations at VG Multimedia.

The Challenge

Before deploying VCS, vg.no used sniffers² to track the traffic on their online publications. They found this to be a resource-intensive, outdated and slow way of monitoring their traffic that also processed a lot of redundant data, as opposed to real-time statistics.

The Solution

Now that vg.no uses VCS to aggregate data, they no longer need to go through millions of logs to detect a certain type of error. "While we were using the sniffers and we had problem with a service, we were unable to figure out what the hit rate for a certain host or a particular URL was. Now VCS helps us to immediately detect an abnormal pattern, view the hit rates to identify the problem and quickly resolve it. If we see something strange is happening, for example if one of our publications has an abnormally low cache hit rate, it's now easier to drill down into the statistics to figure out what is wrong," Audun said.

Varnish Cache, the open-source web accelerator, was originally developed for and in collaboration with vg.no.

A network sniffer monitors the data flowing over computer network links. It can be a self-contained software program, or a hardware device with the appropriate software or firmware programming. Also sometimes called "network probes" or "snoops," sniffers examine network traffic, making a copy of the data, but without redirecting or altering it.

How Does VCS Work?

VCS is a very flexible tool for aggregating statistics and displaying them, based on your own preferred set-up. You can define the grouping of the statistics by using the Varnish Configuration Language (VCL) to set up keys.

This allows you, for example to track statistics by virtual host and domain, by URL, or any other way you see fit, based on your VCL key definition. The data collected can be queried and presented in top lists, according to various parameters.

You can, for example, track the URLs that take the longest time to load, the URLs that get most hits, or the articles and products that attract the most traffic, etc. The possibilities for using VCS are practically endless. At vg.no, VCS has been integrated with Graphite, a tool that visualises the data coming from VCS.

The Benefit

VCS is an outstanding tool for monitoring real-time data. Due to the easy configuration of the VCL, the tool is very flexible and easy to use. "Using our old system we could, at best, view statistics with a 5-minute interval.

With VCS, we can set the interval to 10 seconds or 1 minute, which allows us to react to problems very quickly.

In our industry, speed is imperative and VCS helps us to give our readers a much better user experience," Audun told us.

Current Use

vg.no is currently using VCS for operational debugging and to identify security issues in real time. "I'll notice immediately if one of our publications has an abnormally low hit rate. And using VCS, I can quickly figure out why,

By creating different criteria or keys to view the exact statistics I need to get to the root of the problem," Audun said. "I've defined a library of keys in VCS to log information that is important for me and that helps me during the debugging process."

"VCS is the best thing that has come from Varnish Software since Varnish Cache. It's a straightforward tool that is easy to use and it really helps us to figure out what is happening in our user traffic."

The Future

For the journalists at vg.no it is of utmost importance to know which articles are popular and then turn around and write more about that same topic. To give journalists this insight faster, Audun plans to use VCS to push real-time data to them. "Our journalists check every three minutes how many people have read their articles. With VCS it is possible to update this information every 10 seconds." Audun adds that doing this will greatly enhance user experience and loyalty, which will have a positive effect on the news providers' bottom line.

Audun tells us that another very useful VCS feature that vg.no is planning to use is the ability to track the most popular articles at each moment in time, and push that information to end-users. "With other systems, such as one written in PHP, this is information that is really difficult to keep track of, but with VCS it's easy. I use an article ID as a key.

The article ID counter is then incremented every time someone reads the article, and the information is updated every 10 seconds or so. VCS does this effortlessly while this would require a lot more resources in PHP." Finally, vg.no has plans to put Varnish Cache in front of their advertising system and use VCS to track certain key performance indicators for their various ads. This will allow them to give advertisers better information on how their ads are received on their publications.

Summary

VCS captures and provides real-time aggregated statistics on all traffic that goes through Varnish Cache. It is a flexible and easy-to-use tool that allows web developers to detect abnormal traffic patterns or performance issues in real time and to act quickly to resolve the problem. As a result, any website's performance can be optimised, the user experience enhanced and loyalty increased, significantly improving the bottom line of digital media and other online businesses.

About Varnish Software

Varnish Software is the provider of the open source web acceleration software Varnish Cache. Varnish Cache is a reverse HTTP proxy that will speed up a website by storing a copy of the page served by the web server in its cache the first time a user visits that page. The next time the user requests the same page, Varnish will serve the copy instead of requesting the page from the web server. This means that your web server needs to handle less traffic and your website's performance and scalability go through the roof. In fact, Varnish Cache is often the single most critical piece of software in a web based business.

Varnish Software delivers a comprehensive range of enterprise products and services that help customers further increase the performance and scalability of their websites. Leading websites all over the world rely on Varnish, including BBC, Nokia, The Morningstar, Qt / Digia, The Globe And Mail, The Hindu and Vimeo.

The Varnish open source project began in 2005 as an idea within VG Multimedia, Norway's largest online newspaper and Varnish Software was then founded in 2010. The company has offices in London, Stockholm and Oslo.

Los Angeles - Paris - London
Stockholm - Singapore - Karlstad
Dusseldorf - Oslo - Tokyo



VARNISH
SOFTWARE
www.varnish-software.com