


VARNISH
Makes Websites Fly

Whitepaper

The high cost of poor web performance

What to consider when choosing a caching solution



The Internet has not only given rise to new types of business models and companies such as Facebook and Google, it has also changed the way business is conducted, as the incredible success of Amazon confirms. What is perhaps more important is that the Internet has also changed users' expectations when it comes to interacting with any type of company on the web.

Looking at the most successful Internet companies can inspire others. There are two important factors that all of these companies are mindful of and account for in their websites:

- Internet technologies and business models are constantly changing and need flexible, forward-compatible technologies that are easy to scale and integrate.
- Users expect websites and web applications to be fast and reliable, otherwise they will look for other alternatives that are just a mouse-click away.

The Internet –rooted in open source

The Internet itself has been the best case for Open Source (OS) since Sir Tim Berners-Lee put it into the public domain. Its infrastructure largely depends on OS software and open standards. The most popular content management systems including Wordpress, Drupal and Joomla are OS; one of the most widely used databases, MySQL, is OS and according to Netcraft, 85 percent of all websites are hosted on OS web servers.¹ Additionally, Builtwith.com reported in 2013 that Varnish Cache was the fastest growing web server in 2012, growing by 38% from 2011 to 2012. Varnish Cache is used by more than 500, 000 websites worldwide including the BBC and the New York Times.²

The new tech giants Google, Amazon, Facebook and Twitter are built on OS platforms and only have a proprietary layer on top of it. Even Apple bases many of its platforms and products on OS.³

Therefore it is very likely that many companies are already using OS without actually knowing it.

Why OS?

There are good reasons for Internet companies to use OS technology. Most people think that this is purely for cost reasons as OS solutions range from free to low-cost. However, this is probably not the first reason for Internet companies (though a welcome side effect) but rather the control and the flexibility OS gives them. Having access to the source code allows companies to customize the software as much as they need to. As the Internet business is work-in-progress with business models and technology continuously changing, OS allows them to adapt and fine-tune their websites and applications to gain or sustain a competitive advantage. Furthermore, software that is based on open standards and open protocols is forward-compatible which is extremely important in this environment where technology is progressing fast.

Open standards also make software easier to integrate and scale with other systems, which is very important for fast growing companies. Last but not least, the OS community gives companies access to a huge pool of resources, new ideas and developments.

For more established companies and government organizations, being able to avoid vendor lock-in is extremely important. The more critical the technology, the more important it is for them to keep control over it and not be at the mercy of a vendor.

It is not only the OS layer that makes these internet companies successful. However, OS gives them the flexibility and speed to adapt their products in line with what users want and expect. And this ability to deliver a superior user experience is a key success factor.

1. <http://news.netcraft.com/archives/category/web-server-survey/>

2. <http://trends.builtwith.com/Web-Server/Varnish><http://trends.builtwith.com/Web-Server/Varnish>

3. *Economic Impact of Open Source on Small Business: A Case Study* By Mike Hendrickson, Roger Magoulas, Tim O'Reilly, July 2012

Why performance matters

If there is one common attribute that underpins the success of such diverse Internet companies as Amazon, Twitter and Google, it's performance. Any Internet company from simple business websites to social media sites to e-commerce sites must be reliable and fast – very fast!

A New York Times article quotes Harry Shum, a Microsoft computer scientist who said that “users will visit a Web site less if its loading time is slower than its competitors by 250 milliseconds, or one-quarter of a second. That is less time than a single eye blink.”⁴

According to Research from Aberdeen Group, one second delay in page response will result in 11 percent fewer page views, a 16 percent decrease in customer satisfaction and seven percent reduction in online sales.⁵ In monetary terms this means that if your site earns somewhere around 75,000 Euro every day you could lose nearly two million Euro in yearly sales.

But it is not only the fact that users expect websites to load fast that makes speed a determining success factor for an online business. The perception of time is also changing: it's accelerating. Comparing research data from 2009 with similar research from 2005, Forrester Research measured a significant change in the perception of time. In 2005, online shoppers expected a page to take at least four seconds to load and waited patiently for it to do so. Four years later, in 2009, they expected web pages to load in two seconds or less and would typically abandon a site after three seconds.⁶

One of the reasons for this change in perception is the growing number of competitors: users now have far more options to choose from and thus have a stronger empowerment to abandon a site if it is not performing according to their expectations.

What causes websites to underperform?

There are many reasons why websites are slow. The main reason is increased server load. Since the beginning of the Internet, web servers have had to process more and more objects on sites and cope with larger page sizes as well as growing traffic numbers and cyber-attacks. Meanwhile, the underlying caching architecture designed to handle these loads hasn't changed much.

Caching is a way to serve up web content to documents rapidly. It works by creating a temporary storage area (or cache) to store information so that the computer does not have to keep re-accessing it from the central server. However, most of the caching solutions still follow the old concept whereby a computer has a primary (RAM) and a secondary storage (disk). As such, these caching systems need to track how much RAM and how much disk space they can use and move content back and forth between these two storage areas. This costs time and CPU power.

Varnish Cache was designed from the very beginning to address this problem. It is built on the premise that computer memory is virtual. It stores data in virtual memory, leaving the task of deciding what is stored in RAM and what gets paged out to disk to the operating system. This saves CPU power and time and is therefore much more efficient in how it uses storage.

4. http://www.nytimes.com/2012/03/01/technology/impatient-web-users-flee-slow-loading-sites.html?_r=2&
5. <http://www.aberdeen.com/Aberdeen-Library/5136/RA-performance-web-application.aspx>
6. http://www.damcogroup.com/white-papers/ecommerce_website_perf_wp.pdf

Reverse what?

Varnish is also different from other caching solutions because it was designed as a 'reverse proxy'. A reverse proxy acts just like an ordinary web server in that it intercepts all web requests before they reach a company's web server.

The advantages of a reverse proxy are that it caches the requested pages from the server, so the server avoids having to reproduce multiple impressions of the same page. This means that the next user can view the page without experiencing a delay. By doing so, it is able to serve hundreds of consecutive requests per second, speeding up the website performance and reducing server load drastically.

Fighting the high cost of underperformance and threats

The costs of underperforming websites and cyber-attacks are high for online businesses today.

Downtime alone can result in:

- Lost employee productivity costs
- System restoration costs
- Lost sales opportunities
- Lost customer costs
- Damaged reputation costs

There are several reasons for downtimes: from web traffic peaks that overload servers to application problems to cyber-attacks. The latter is becoming a growing problem. According to the UK Cabinet Office, in 2011, cyber crime was estimated to cost the UK about £27bn a year, £21bn of those costs were business related. But these are just estimates:

"In all probability, and in line with worst-case scenarios, the real impact of cyber crime is likely to be much greater."⁷

Varnish cache grew in popularity when it helped to minimize the impact of one of the biggest DDOS attacks in history. Wikileaks used Varnish to deflect an attack that exceeded 10 gigabits a second, demonstrating that the ability to handle high traffic loads makes the caching solution an important tool to manage Internet security.

Even if your company is not vulnerable to attacks or other causes of downtime, it still needs to be ultra-fast for the competitive reasons outlined earlier. This is why content-heavy, dynamic websites such as the BBC, Facebook and Twitter were among the first users of Varnish Cache.

Should I use the OS version of Varnish or invest in a Varnish Subscription version?

Varnish Cache is available in two different versions depending on your requirements:

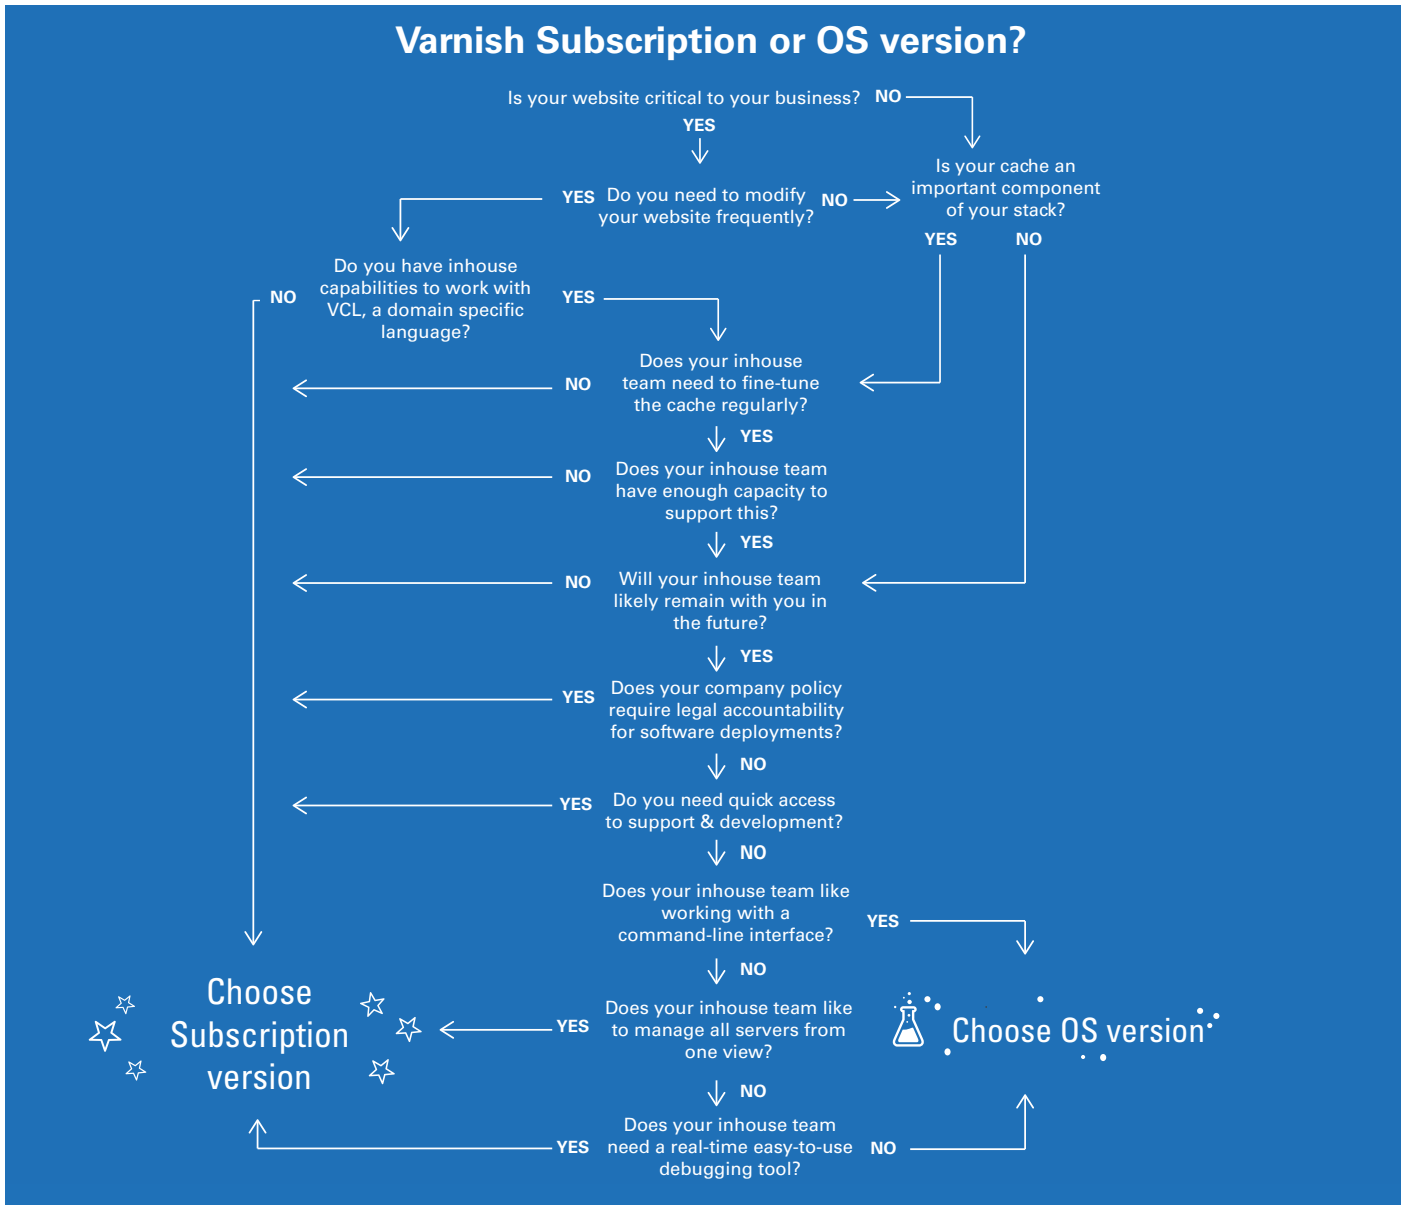
- An OS version, with free license under a two-clause BSD license (also known as the FreeBSD license)
- A subscription version that includes support and proprietary products and dl solutions that bring additional value to the OS version of Varnish Cache.

Although it may be tempting to opt for the OS edition, companies need to consider the **total cost of ownership** over the long-term, factoring in their requirements for support and customization and legal considerations. Some enterprises, usually larger ones, have more development resources in house to support these, but others may find better overall value in investing in an enterprise subscription. The diagram below can help you determine whether the OS or subscription version of Varnish is most appropriate for you and your organization.

Varnish Software offers support, training and consultancy to a wide range of customers of all sizes. The insight into customers' business problems further supports the development of the community edition as well as added components available with an enterprise subscription.

7. <https://www.gov.uk/government/publications/the-cost-of-cyber-crime-joint-government-and-industry-report>

Varnish Subscription or OS version?



The business benefits of a Varnish enterprise subscription

As described above, the websites that benefit the most from a Varnish subscription and an optimized version of Varnish Cache are the content-heavy, dynamic sites that have a large number of concurrent users and spikes in traffic. Wetpaint.com is an example of a Varnish customer that signed up for a subscription and with help from the Varnish support team saw a significant improvement in their web performance. In November 2012 they implemented an optimized version of Varnish Cache and their site went from being unresponsive and slow to becoming very fast. According to Tony Flint, Wetpaint's Senior QA, IT and Operations Manager, Wetpaint's average load time went from 8-9 seconds to less than 3 seconds. And after setting up Varnish they were also able to cut the number of servers supporting their site from 10 to 4. "With Varnish we were able to cut our hosting related costs by 50% and we had a much faster website," Flint told Varnish Software in an interview.

One of Varnish Software's very first customers, Norwegian online newspaper Vg.no, was experiencing challenges with scalability and speed. Before implementing an optimized Varnish installation their objective was to make their website faster and to save costs on their web infrastructure. They attained their goal and managed to reduce their hardware costs by 90% and they saw a dramatic reduction in response times.

For those companies with the mission-critical performance and up-time requirements outlined in this paper, the enterprise subscription includes a number of additional services (depending on the subscription plan you choose) including round-the-clock support, open source assurance and software maintenance and proprietary products and solutions such as the Varnish Administration Console (VAC), Varnish Custom Statistics (VCS) and Varnish Paywall solution.

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 will 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.