The Truth About WordPress Performance

Why You May Not Need What You're Being Sold

copyblogger

Table of Contents

The Content Publisher's Mission

If you're publishing online content – either for marketing or with an ad-supported business model – your mission is simple, but not easy. Get as many relevant site visitors to experience your education, entertainment, or news.

It's not easy because you manage a complex digital orchestration. From content creation to design and everything in between, you need to create an experience that both site visitors and search engines engage with in a positive manner.

The experience needs to be compelling and it needs to be delivered quickly. It needs to, in a word, *perform*.

Performance – The New Problem Space

Performance is one of the biggest challenges that site owners face. Pages need to load quickly because attention is quickly lost.

The human brain is trained to the "performance" of a printed page. According to Google, when a reader turns a page in a book, it takes the mind just 100 milliseconds to process. Not coincidentally, that is Google's goal for web performance.

Right now, for online content, the threshold for poor performance is just two seconds, or 20 times what it takes to view a printed page. And this gap of tolerance is quickly shrinking.

Google considers a web page "slow" if its load time is over 1.5 seconds. A study published by a University of Massachusetts professor found that people are "pretty patient for up to two seconds." At just 3 seconds, 57% of them will abandon the page.

Basically, if your content does not load quickly, your readers are not going to stick around. And with search engines factoring page load times and bounce

copyblogger

rates more heavily into their ranking algorithms, the negative impact of poor performance is exponential.

You're likely an expert in your field, but you are probably not an expert in site performance tuning. Over the last few years, an entire industry has popped up promising to provide you with worry-free, one-click solutions to performance issues. While there are some good, affordable solutions that aid in site performance, there is a lot of snake oil out there as well.

Still, a website that performs well is achievable and affordable. We've written this whitepaper to help demystify the topic for you and give you the tools you'll need to pick the right solution for your specific performance needs.

Why listen to us? At Copyblogger Media, we are content publishers just like you. We've been in your shoes. We *are* in your shoes. Every day.

And we have a big advantage: our Synthesis team and our business partners at W3 EDGE really do understand the topic of performance. So we challenged them to write a simple guide that that will help you pick the right performance solutions while avoiding the path of false promises.

WordPress – A Tale of Two Users

WordPress is an open source content management system ("CMS") written in a series of scripting languages that are processed either by the server or an end user's browser.

As the most popular CMS online, WordPress has revolutionized online publishing by allowing the writer to create content without having to involve an HTML developer every step of the way. For this flexibility, however, there is a resource price to pay.

Writing   Reading

The key to WordPress performance is realizing that the only users who need to access the site in resource-intensive ways are the content publishers. On the other hand, readers of the content can access an optimized version of the content and still receive the full experience.

The Performance Ecosystem

If one or more companies in the site performance ecosystem haven't bombarded you with marketing already, consider yourself a statistical outlier.

The performance, hosting, and acceleration of websites is a space that has really heated up in recent years. It's to the point where it is actually growing crowded.

The ecosystem consists of hosting providers, "cloud hosting" providers, page acceleration caching services, CDN (Content Delivery Network) providers, professional tuning services, and managed hosting providers that combine one or more of the above.

In the remainder of this whitepaper, we will define each of these provider types and clarify how and when they can help you. Moreover, we will direct you to simple online tools that can help you determine what you may (or, more likely, may not) need based off quantifiable data as opposed to hype.

Origin vs. Edge

One core concept to understand is the difference between the origin and the edge when it comes to optimization.

The origin is the source of your content—your WordPress installation. The edge refers to services removed from your actual origin WP installation, and ostensibly closer to end users, that might distribute or perform optimization on your content. Edge services include Content Delivery Networks and acceleration services.

To illustrate, consider the difference between origin and edge as being like making a purchase at Amazon.

Some orders are fulfilled by from the sellers' actual warehouses. This is the equivalent to origin delivery. Other orders are fulfilled by Amazon from a

warehouse that is closest to the customer purchasing it. This is the equivalent to edge delivery.

The edge delivery sounds more efficient, right? Yet, if Amazon is out of stock of the manufacturer's product, there is latency. The manufacturer must send it to Amazon, where it then has to be received, checked in, and made ready for shipment before Amazon can fulfill the order. Each step is time-consuming and could have been avoided if the manufacturer would have simply drop-shipped the product directly.

A key takeaway here is that the manufacturer let the product get out of stock at Amazon in the first place. This indicates an un-optimized origin. Even in Amazon's world, an un-optimized origin trickles to the edge and negatively impacts customers. The same applies to WordPress.

An optimized WordPress origin can be further enhanced with edge services. But trying to disguise origin issues with edge services can often lead to more issues rather than solving any problems. Tuck this understanding away, as we'll talk more about this when we look at origin vs. edge caching.

Hosting and Cloud Hosting

Hosting is a core origin technology. And there is no question that good hosting is a fundamental building block for a well-performing site. But the introduction of the buzz term "the cloud" has really made it hard to distinguish good hosting from good marketing.

Good hosting providers combine top-tier hardware running optimized software with great data center network connectivity. How do you figure out whether a hosting provider falls into this camp? There is no way to tell for sure, but a good sign is the disclosure of specific details. If a hosting provider does not have any information about their data centers, then you should ask them … and they should tell you.

Synthesis, for example, specifically leverages the hardware and data center connectivity of a specific provider with whom we have a commercial contract. It is one of the best data centers in the country, and we'll tell you all about it. We will also proudly tell you about the data center facilities of other Synthesis components, like our managed backup infrastructure located in Dallas, TX.

Network Cabling in Dallas, TX Synthesis Data Backup Facility

Incessant marketing messages may have you wondering: with the onset of "the cloud," doesn't all of this change? This buzz term is actually quite misleading in most hosting-related marketing.

"The cloud" is a broad term that alludes to the notion that you will be able to access thousands of scale-out servers, network connections, and caching schemes, which, when combined, will allow your site to scale from one visitor to a billion visitors without issue. This is misleading, to say the least.

copyblogger

The truth is that cloud technologies really don't benefit the end user nearly as much as they are marketed. They do, however, benefit IT departments and hosting providers by offering layers of flexibility that lower operational costs.

For example, at Synthesis, we use cloud technologies to quickly resize customer servers and move them from one hardware platform to another. We also use cloud technologies to aggregate millions of server log entries and look for site health and security patterns. This helps us, but it doesn't benefit site owners by making their sites faster.

What makes for a good hosting platform? Good hardware, finely tuned software, and high-speed bandwidth.

But there is more to having a high performance site than just good hosting. A well-optimized origin is the key, and this is where origin caching comes into play.

Caching

Caching is a tried-and-true performance technique that has been in use since the beginning of the modern computing era. To best explain, we will compare a published web page to … Scandinavian furniture.

When you go to a Scandinavian furniture store, the showroom is full of sleek, contemporary furniture. Yet, if you've ever purchased one of these pieces, you

know that you get a box of parts that has to be assembled. This process takes time and leaves you asking, "Why couldn't they just sell it to me assembled?" The reason is that shipping and storing assembled furniture is expensive. It takes up a lot of space.

The same applies to content management systems like WordPress.

WordPress stores content efficiently as text in its database and then combines the functionality of plugins with the design elements of a theme to make the finished page on a site. Though not as cumbersome as the assembly of Scandinavian furniture, it is a resource-intensive process and does take time. Like the assembled furniture you see on the showroom floor, caching is the act of storing a web page in its final form so that it loads quickly.

The WordPress core was built with the concept of caching in mind. Specifically, it has the ability to support a variety of cache implementations on top of the core content publishing functionality. This means that caching implementations are treated as first-class citizens versus being treated as an afterthought that is poised to break as the WP core evolves.

The WP community did this for two reasons:

1. They understand the difference between the resources needed to serve static pages to the visitors of a website and the resources needed for content creators and the publishing features they demand.

2. The community realized the importance of implementing caching and other optimizations as close to the origin of the content as possible.

Origin Caching vs. Edge Caching

Though caching can come in a variety of flavors, it is fairly safe to boil this conversation down to two forms of caching—origin and edge.

Origin caching is a flexible, stable form of caching that happens at your site. In the same way that your Apple laptop caches data in memory to optimize the user experience, components can be added to WordPress to help it deliver assembled furniture, so to speak.

We believe in caching as close to the WordPress core, or origin, as possible. Data presented in DashboardJunkie.Com's "Test of WordPress Caching Plugins" shows the profound positive impact caching has on key performance and scalability metrics applicable to WordPress sites.

Like the article's conclusion, our preferred method of doing this is by using the WordPress caching framework developed by W3 EDGE called W3 Total Cache. This framework allows you to specifically tailor caching for your site's needs, and it can serve as a communications conduit to other elements in the performance ecosystem such as front-end caching services and CDNs.
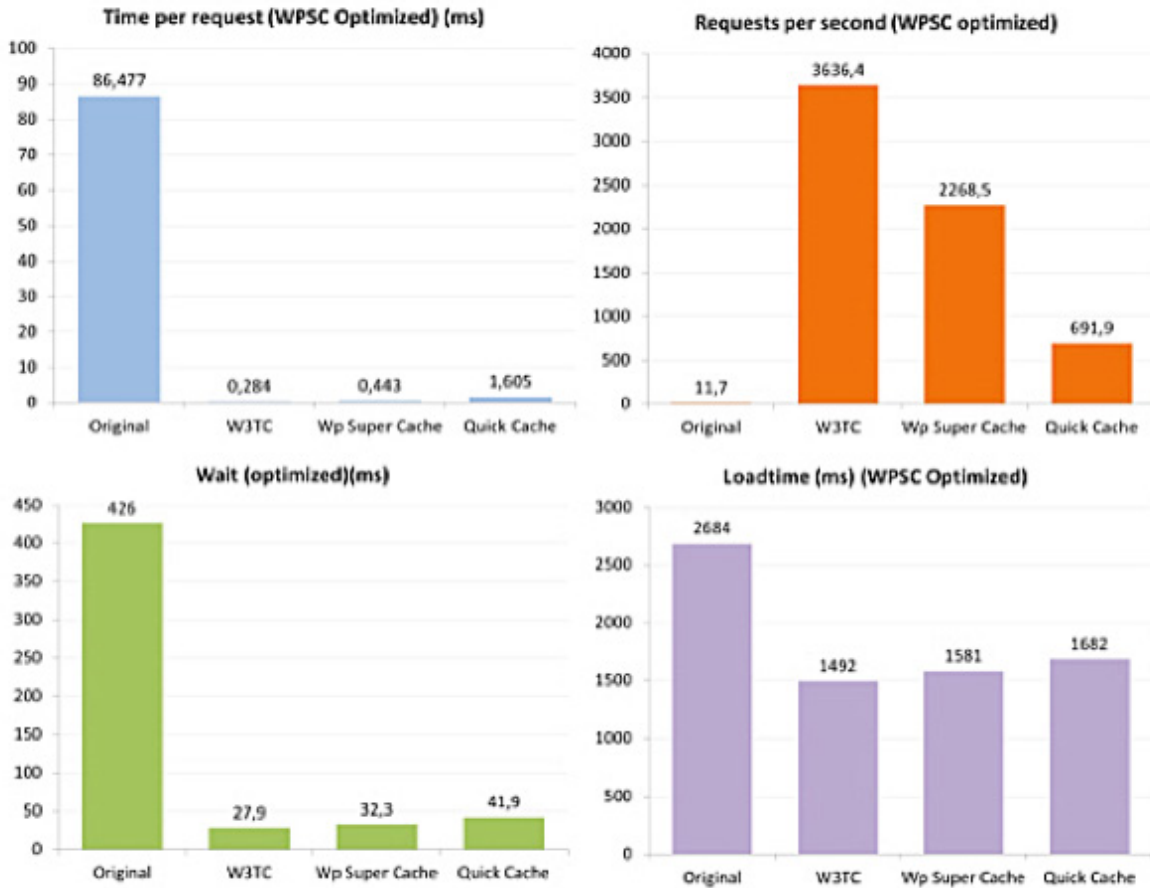
Test of WordPress Caching Plugins by Kim Teztlaf Courtesy of DashboardJunkie.Com

Edge caching is another method for caching websites that is typically used in very large website platforms. This method uses one or more dedicated servers in front of a website's origin to cache the assembled pages and serve them to visitors. These edge servers could be physically located very close to the origin server or they could be quite far away.

copyblogger

Examples of edge caching services include CDNs and full-page cache accelerators like Akamai. Though originally introduced in the enterprise segment, new offerings are beginning to enter the marketplace with commoditized pricing and are more accessible to small business owners.

While edge caching can speed up website performance, these services are designed to accommodate millions of customers and are not necessarily tuned for an individual website's needs. For example, your site may need granular minificiation tuning that includes some files and excludes others, but this functionality is unlikely to be available from a page acceleration service.

Moreover, if edge caching is implemented without a solid origin caching strategy, it often leads to okay but not great results. (This is exemplified above in our Amazon analogy.)

Good edge cache services replicate cached content to multiple edge destinations. For example, a UK-based website could cache static objects or entire pages to edge servers in both Japan and Brazil. Consumers in Japan would take delivery of cached pages from the Japanese edge servers and Brazilian consumers would take delivery from a Brazilian-based edge server. The result is that edge servers can help by getting your cached content physically closer to prospective readers who are far away from you, which allows the content to load faster.

There is a catch, however …

The Telecommunications Reality

Not all consumers of content are near an efficient edge.

At Synthesis, for example, we've had customers from Latin America inquire about using a CDN in order to more efficiently service their Uruguayan-based readers. The problem is that the service they selected did not have edge services in Uruguay. Thus, their customers were hitting edge servers in Miami.

According to telecom market research firm TeleGeography's 2012 Internet connectivity data, 84% of the data connectivity in Latin America flows out of region, with 58% of it hitting a single city—Miami. Furthermore, 79% of Latin America is geographically located east of the United States. Thus, Miami is a good place for an edge server.

So … if your origin server is right up the telecommunications coast in Virginia, is there significant benefit to providing an edge in Miami? In our real world experience with the above scenario, the answer was no.

2012 Global Internet Map Courtesy of TeleGeography (www.telegeography.com)

On the other hand, Europe's connectivity is the exact opposite of Latin America's. According to TeleGeography, 70% of Europe's bandwidth is routed among cities in-region. Thus, connectivity between cities is very rich. A simple edge server strategy in London, the world city with the highest amount of network bandwidth, could theoretically suffice for all other major cities in region.

These facts reveal the real problem: tertiary markets.

Our website StudioPress.com, hosted in Virginia, has a sub-second full-page load time from Dallas, TX. If we travel 99 miles away to the East Texas town of Mount Vernon, TX, our load times start to creep up to the 1.00 second mark. Why? The consumer Internet connectivity in this outlying region isn't what it is in Dallas. Moreover, there is no CDN or acceleration service with edge servers in this region—the closest ones are in Dallas and still require a trip down the problematic telecommunications leg.

What does this seemingly complex telecommunications talk mean? Simply that the consumers of the content on our websites come from everywhere, and each has a telecommunications travel story that is different.

If you want to perform both domestically and globally, you have to cater to worst-case scenarios. This means it is imperative to optimize content at the origin *and* host the origin content on quality infrastructure. Surrounding poorly optimized content and/or poor infrastructure with edge services does not guarantee success.

Content Delivery Networks and Full Page Caching

In the previous section, we talked in general terms about edge caching. Let's get more specific.

There are two popular forms of edge caching services being marketed today. The most established are Content Delivery Networks(CDN). CDNs are designed to cache and deliver static objects through distributed edge servers that are located in-region. Static objects are website components such as images (e.g. JPG or PNG files), style sheets, audio files, or even videos. Typically, CDNs do not deliver fully cached pages.

Let's put this in the context of our furniture analogy from earlier.

Imagine buying a contemporary nightstand online. Since the manufacturer packages all of the components in a small box, it is somewhat efficient for them to ship it to you. However, this particular unit features a very heavy and bulky glass top.

If you could take your receipt to their local store in your city, it would be more efficient for you to pick the bulky glass top up locally versus having to pay for shipping the heavy object. A CDN works in the same way: users get inefficient parts of a web page from a local source, but they do not get the fully assembled nightstand delivered to their front door.

Full-page caching services are relatively new on the scene. Whereas a CDN only caches static objects from your website, full-page caching stores fully cached pages from your website and delivers them in-region from a vast network of edge services. These services have put together some impressive technologies to accomplish their goals.

copyblogger

As impressive as these technologies may be, we posit that usage of either a CDN or a full-page caching service on a site with an un-optimized origin is not going to produce the results that you want. But the usage of one of these services with an optimized origin? *That* can produce amazing results. Just remember what needs to come first.

Managed Hosting Providers

Managed hosting providers are a new breed of hosting provider, offering advanced service levels for performance, support, and security. From a performance perspective, many managed hosting providers provide free edge caching services as a part of their service. These can come both in the form of CDNs and full-page caching services.

Until now, we've only alluded to the fact that we don't believe in one-size-fits-all performance solutions. Let's go ahead and make it clear—***we don't like one-size-fits-all performance solutions.***

Most managed hosting providers will improve your site's performance with their system-wide edge caching and optimization techniques. This will certainly put your site on a different playing field than a commodity host. However, why just be good when being great is affordable?

At Synthesis we believe in providing great hosting components (hardware, software, and connectivity) combined with optimal *origin* caching services.

Why? We believe in optimizing the origin before considering more generic edge-caching services. We believe that edge-caching decisions should be made based on empirical data that is gathered *after* origin optimizations have been made.

Professional Tuning Services

If you have ever ridden in a BMW, you know these cars can be very fast. Aftermarket parts can make a BMW faster, but they can also void your warranty. The same applies to letting your developer do something like modify the WordPress core—it's just not a good idea.

In the BMW world, companies like Dinan moved in to not only make parts but also to provide tuning services to make BMWs even faster. Best of all, Dinan did so in a way that was compatible with BMW warrantees, creating a win/win situation for the customer. The same concept exists in the WordPress world.

One of the most recognizable such "part makers" and tuners is W3 EDGE, the creator of the free plugin W3 Total Cache, which is the antithesis of performance "fairy dust." W3 Total Cache is a true framework that sits between WordPress and all the major performance tools and services that can be applied to your site. As a framework, it can be configured to meet just about any site owner's needs.

W3 Total Cache supports new high-performance web servers such as NGINX and LightSpeed, as well as front-end caching services like CloudFlare. It also provides support for most major CDNs. W3 EDGE, provides a variety of services ranging from simple caching configurations to full-fledged server tuning.

Solve For Real Problems With Real Data

We all have the same weapon available to us in the battle for better performance: data.

Performance data is accessible and decipherable. You do not have be a site performance expert to become an educated consumer who will select the right performance solutions for your site. And by avoiding the wrong solution, you will save time and money.

You have probably read or heard something like this numerous times: *"Use a CDN. It will make your site faster!"* CDNs can indeed alleviate specific performance issues, but they should not be used with the expectation that they will be a magical potion that solves everything.

Why? Follow the data and you will find your answer.

Data & Cost-Benefit Analysis

WWe have been clear in this document that we believe origin optimizations should be your first priority. The data backs this assertion up.

In the piece referenced earlier [that compares WordPress caching plugins](#), author Kim Tetzlaff provides empirical data showing that using W3 Total Cache to provide simple page caching-based origin optimizations to a WordPress site results in 300X improvement in terms of requests per second that the host server can process. These are not the types of optimizations that you want to try and push out to edge caching services.

Still, before you throw any type of caching solution at your site, your first task as a site owner is to understand whether your site has a performance problem. You do this by looking at the facts versus getting caught up in a hyped-up, buzzwordy "grade" that a generic page checker gives you. There are a number of tools available that use very technical factors to give you such a grade. These are misleading and often don't provide you with data points you can act upon. We don't use them.

The fact is that websites load as a series of objects. Much like the contents of a box of Scandinavian furniture, these objects all come together for the final product. There are several simple, free online tools that can help you to look at these objects and figure out where the performance problems are. Unlike all of the dowels and connectors that come in a furniture box, these analyzers actually show you the data in a simple and meaningful way.

Here are two:

1. http://tools.pingdom.com/fpt/

2. http://www.webpagetest.org

Let's run http://websynthesis.com through both of these tools.

Using Pingdom first, the Synthesis website loads in 809ms, or *less than 1 second*. From an end user perspective, and from a search engine perspective, this is great.

However, in the Web Page Test report (#2 above), you will see a few grade marks of "F". An F is pretty harsh for a sub-1.00 second site. Why the failing grade marks?

One failing mark is because we don't compress images … which makes this a good time to define an important rule: ***act on data, not on emotion***.

Let's see if these images whose non-compression apparently warrants a failing grade are actually a problem. (Either page check tool can be used for this.) The image "icon-performanc.gif" loads in 11ms. This is fast! If we compress this image and get it out the door 50% faster, what's the return on our investment? 5.5ms? Do you think your readers will notice? Another fact is that the browser will cache this image, so it disappears from the mix the second time the page is loaded.

Thus, this "failing" mark from above is misleading!

To get to the facts, change the Pingdom tool to "Sort By Load Time" and you'll see why the site loads in 809ms as opposed to, say, 400ms. One is a script delivered by stats.wordpress.com and another is Google Analytics. We also have to wait on Gravatar images to be delivered for our customer testimonial section. The scripts and images are being loaded from other servers. Of course, we can't control their servers. No front-end caching service or managed host is going to help out here.

The only feasible way we can improve the performance of the Synthesis website is to remove the scripts or scrape the Gravatar images to local files, neither of which we have chosen to do. The scripts are well worth the slight addition to load time because of the benefits they provide. The images are too.

This type of cost-benefit analysis for page load speed is an essential exercise for any site owner committed to delivering the best performance possible to readers. As the saying goes, "you can't put lipstick on a pig." In this context, it means that a site bloated with codes and scripts, especially external ones, can only load so fast … and there is only so much that can be done to speed it up.

There simply is no better performance enhancement than loading only what you and your readers *need* to have on a webpage. Once you've determined that, then optimize it.

You can't dump everything but the kitchen sink onto your pages and then expect handfuls of performance-enhancing fairy dust to mask your bloat and deliver sub-second page load speeds.

We host websites and we are content marketers at heart. Thus, we love data. And data as described above helps us to make informed decisions—such as knowing when we can blow off "failing" grade marks and realize that our site's sub-1.0 second load time is great!

Everything you need to know about your website can be seen with a simple web analyzer.

"Big Numbers"

In many cases you will not even need to look at the objects of a page to find out where your problems lie.

There are two numbers at the top of the Pingdom Full Page Load tool that are very telling—requests and page size. If your home page has more than 150 requests or the page size is greater than 2MB, you have a simple problem: too much going on. (We actually don't like to see more than 125 requests and 1.2MB, but we know that numbers slightly higher than this are tolerable.)

And when your site has too much going on, it's time for a cost-benefit analysis as described in the previous section.

If your site has "big numbers" there are a number of places you can look to reduce your page load bloat:

- How many posts you are displaying?

- How long is your page? (You do realize, right, that most people won't scroll very far?)

- How many large images are you loading for which compression *could* make a tangible difference?

One very helpful tool in paring back a website is Google Analytics' In-Page Analysis. This tool provides a small heat map visualization that will tell you where people click on your page. In most cases, if you have a home page that requires 10 scrolls of a mouse to get to the bottom, then Google In-Page Analytics will show that fewer than 2% of your visitors are willing to go more than one to two scrolls. Translation? Your readers aren't reading it … so get rid of it and pare that home page down!

Another cause of big numbers for a page is poorly built themes or plugins (e.g. social sharing plugins) that place scripts and images inefficiently on your site. Clean themes and the judicious use of plugins are key.

If you have a site with "big numbers," the single most effective solution is to fix it. Keep what you need and jettison the rest. Reduce the big numbers! Beat back the bloat!

As an alternative, you *could* consider a distributed front-end caching service. This can help to offload your problem onto their network. Just remember that this method only treats the symptoms, not the disease. But at least if you have data, you can make an informed decision with your eyes wide open rather than simply acquiescing blindly to marketing hype.

Taking Action

Once you have gathered data, it is time to take a few steps toward improving your site's performance.

The biggest mistake that site owners make is rushing to use advanced solution components before taking care of the basics. The mistake is the same as putting a high performance air filter on a car with a small engine designed for fuel efficiency — the engine is the core of the performance.

Efficient Foundation

At Synthesis, we routinely see sites that suffer from issues related to "big numbers." These issues include, but are not limited to:

- Bloated CSS and JavaScript

- Inefficient image resizing or slider scripts

- Plugins that simply consume too much memory

Though themes can be reworked to reduce bloat, and poorly designed plugins can be substituted for leaner ones, there are other alternatives for improving performance. None is more effective than implementing an efficient WordPress theme framework.

WordPress' web-wide success as a content management system is due in large part to its architecture, which allows themes to be broken into a base functionality component – the framework – and a design component, referred to as a child theme, which sits on top of the framework. You can update the framework (think of the engine of a car) without affecting the child theme's design (the car's paint job).

There are several good frameworks available today. We are going to focus on the efficiencies that our own Genesis Framework brings to the table.

The Genesis 2.0 Framework by Copyblogger Media makes both performance and security a priority. Genesis provides a solid foundation for theme designers by supporting HTML5 markup and mobile-responsiveness out of the box, thus alleviating the need for such features to be hard-coded. Moreover, Genesis provides key functionality such as schema.org markup and SEO meta tag support, plus easy implementation of Google Authorship, which alleviates the need for several common plugins.

Additionally, the Genesis Framework:

- Utilizes native WordPress APIs to ensure efficient, secure, and compatible PHP execution.

- Implements a clean CSS structure with clearly defined customization files to reduce overhead.

- Can be extended safely through officially supported Genesis-specific plugins that are maintained by experienced community developers.

- Has a W3 Total Cache extension that supports fragment caching, thus ensuring the highest performance optimization available for pages with transient data-driven sidebars, header, and footer areas.

If you are designing a new site or serious about core optimization, an efficient theme framework like Genesis will not only help your designer/developer be more productive, it will decrease the likelihood of "big numbers" creeping into your WordPress site.

Whether you run the Genesis framework, another WordPress theme framework, or have a custom-built theme, you need to understand your theme's basic components so that you can better analyze your site's performance. Three basics to look at are: static objects, dynamic calls, and 3rd-party scripts.

Static Objects

Static objects include images (jpg, gif, png), style sheets (CSS files), JavaScripts (js files), and cached pages created by a caching plugin like W3 Total Cache. A good web server with effective WordPress caching should load these quickly. If your server does not, then you need better hosting.

You could push these out to a CDN or full-page caching service, but that money could be spent on building a better foundation, a more optimized origin, which we've already explained is vital.

Air fresheners will only go so far before you have to simply get the carpets steam cleaned.

Dynamic Calls

WordPress is primarily built in PHP, which is also known as server-side

scripting. Every dynamic call that makes it back to PHP will be much slower than a cached or static object.

Consider this URL:

http://mydomain.com/thumb.php?src=myimage.jpg&w=250&h=400

If a URL looks like this (with ".php?" and then a string of characters after it), the site is likely using server-side scripts to render content. In the URL above, the call is using a dynamic image resizer, which is slow and very costly in terms of server resources. Messy URLs are usually dynamic. You'll need either a theme developer or a site tuner like W3 EDGE to help you eliminate or cache them.

3rd-Party Scripts

The leading cause of slow-loading sites is scripts that are pulled from someone else's network. These include advertising scripts, website analytics scripts, and social sharing scripts that come from Facebook and Twitter.

These offsite scripts are hosted on very powerful networks and CDNs but are being accessed by literally billions of users, and they tend to load slower than objects delivered from your site. The critical point to realize about 3rd-party scripts is that they too can be dependent on remote scripts of their own, creating a "4th-party" script effect.

There is NO magical potion you can throw at 3rd-party scripts. If you suffer from slowly loading external scripts and sign up for a front-end caching service, CDN, or managed host to try and speed up your site, you are likely wasting your money. The only solution is a good ol' cost-benefit analysis. If the benefits of the 3rd-party scripts outweigh their cost on your load time, keep them. If not, bid them *adieu*.

Why Not Try It?

Perhaps you are thinking, *why don't I just try some of these performance enhancers and see what works?* Because there are certain actions you can take with the objective of improving performance that can actually turn out to be costly mistakes.

For example, if you sign up for a CDN just to "check it out," you are likely going to have to change the URLs on your image paths. If you are the #1 hit in Google Image Search for "Labrador retrievers" with your current URL (http://mydogsite.com/cute_lab.jpg) and switch to a CDN, your URL will likely change (to something like http://cdn.mydogsite.com/cute_lab.jpg). This is a different URL, as the subdomain is treated separately from the main domain, and you are now in 301 (redirect) territory with Google. Be careful as you can kill your traffic this way.

Side note: another benefit of W3 Total Cache? It actually supports the "rel canonical" header, which will allow you to better communicate with a search engine when implementing a CDN.

Minification, the process of removing all whitespace from a file, is another topic that people often hop into based on hype. Here is how it works:

Consider the sentence "I like fast, stable websites. I also like mobile responsive websites."

This could be minified to: "Ilikefast,stablewebsites. Ialsolikemobielresponsivewebsites."

You see how the second one is shorter than the other?

Web browsers can actually deal with this in terms of JavaScript and CSS files. There are other benefits like combining files to decrease the overall number of object calls on a page. Still, you are physically changing files … and when things go wrong, they go really wrong, like your site loading nothing but a white screen.

We like minification in theory, and even in practice, but only if there is a data-verified return on the investment. Saving 50% on a file that is only 50 bytes is a waste of effort. Saving 50% on a file that is 300K is something worth talking about.

Granted, the above example is a very simplistic overview of minification. A good minification strategy can provide other benefits such as reducing the number of HTTP calls on a page, which can help reduce "big numbers" and the performance issues associated with them.

The takeaway here is that you should not just mindlessly click minification "On" for your site. Have a strategy, consult your performance tuner, and understand what is actually going on when objects on your site are minified.

Next Steps?

We've covered a lot already, and there is so much more that could be discussed on the topic of performance tuning. But the time has come now for next steps, for actionable takeaways, that you can implement immediately.

First, analyze your site and streamline it. Keep those objects that you need, consider with a skeptical eye those that you just want, and ditch what nobody uses nor even sees uses.

Second, optimize the origin. You do this by addressing the basics of good hosting and implementing origin caching.

- Since most managed WordPress hosts offer money-back guarantees, you can evaluate good hosting with no risk. Switching from generic shared hosting to an optimized, managed, WP-specific host is often the best "performance tuner" possible.

- [Download W3 Total Cache](#) and get basic origin caching in place on your site.

After this is complete, look at your full-page load time. Do you beat the 2.0-second mark? If not, or if you want to get into sub-1.0 second territory, look again at the data. And if analyzing web page load times isn't your favorite thing to do, ask your host for input or consider engaging a WordPress tuner such as W3 Edge.

What Do We Use?

WWe are content publishers. Our primary product websites (copyblogger. com, studiopress.com, websynthesis.com, getpremise.com, and scribecontent. com) run on our own great hosting and use W3 Total Cache for origin caching. We do not use a CDN on four of these five sites. Copyblogger.com uses a CDN because we create infographics from time to time that go viral.

W3 Edge uses a CDN on their primary site in conjunction with W3 Total Cache and good hosting. W3 Edge does not host with Copyblogger Media, but you can rest assured that we are all on the same page in terms of our approaches to performance.

Summary

Ten years ago, good hosting and complex content delivery services such as CDNs where not affordable for individuals and small- to medium-sized businesses. We are very fortunate in today's world that they are. Still, we encourage you to build your WordPress site on a strong foundation, and then use data to evaluate additional performance strategies.

Web site performance matters, and it is affordable and achievable for all. We just want you to know what you're buying and why, because the unfortunate truth is that you may not need what you're being sold.

So ignore the hype, look at the data, and focus first and foremost on the fundamentals of performance – optimizing the origin – before you look anywhere else.

Contributors

About Copyblogger Media

Copyblogger Media is an acknowledged leader in software tools and training programs for online publishers and content marketers. The company powers more than 2 million WordPress websites with its market-leading Genesis framework, the centerpiece of a suite of products that includes StudioPress themes, Scribe content marketing software, Premise conversion optimization software, and Synthesis WordPress hosting. Copyblogger helps marketers maximize these tools with specialized training programs; over 191,000 subscribers receive daily educational updates from Copyblogger.

About W3 EDGE

W3 EDGE began its life as an agency providing marketing solutions across media. Over the past 10 years capabilities in building and optimizing web sites were developed as the organization grew, including: social media, search engine, web performance, conversion rate et al. Along the way it became clear that WordPress was the CMS providing the most value for publishers today and W3 began to focus on creating software that moved publishers forward;

the most well known product along those lines is W3 Total Cache, which is the leading Web Performance Optimization Framework with more than 2 Million downloads and accelerates at least as many web sites.

W3 has either been the technology and / or creative partner for many brands you know like: ASOS, Adorama, AIGA, Edelman, Brian Solis, Center for Disease Control, Constant Contact, CVS Pharmacy, Envato, Hyatt, Kodak, Lord & Taylor, Mashable, Microsoft, Neil Patel, Sanyo, Sherwin Williams, Smashing Magazine, Southwest Airlines, Staples, Sony, R.E.I., Weight Watchers, Yahoo, Yoast and more.

Hyperlink Source List

This Is Your Brain On A Slow Website by Strangeloop Networks
http://www.strangeloopnetworks.com/assets/images/Infographics/Web-Stress-Infographic-500.jpg

Speeding Up Your Website's Database by Paul Tero (Smashing Magazine)
http://coding.smashingmagazine.com/2011/03/23/speeding-up-your-websites-database/

Network Performance: Does It Really Matter To Users And By How Much? By Ramesh K. Sitaraman (University of Massachusetts, Amherst and Akamai Technologies Inc.)
http://people.cs.umass.edu/~ramesh/Site/PUBLICATIONS_files/FinalConference.pdf

Test of WordPress Caching Plugins by Kim Tetzlaff
http://www.dashboardjunkie.com/test-of-wp-caching-plugins-w3-total-cache-vs-wp-super-cache-vs-quick-cache

Global Internet Map 2012 by TeleGeography
http://global-internet-map-2012.telegeography.com

What you don't know about "fourth-party calls" could be hurting your users by Grant Ellis
http://www.strangeloopnetworks.com/blog/what-you-don-t-know-about-fourth-party-calls-could-be-hurting-your-users/