

Web Performance ROI



Content

- A brief history of web performance ROI 3
 - 2010: Should we care about page slowdowns as much as we care about downtime? 3
 - 2011: Can we compare apples to apples? 4
 - 2012: What can we do with a LOT of user data? 6
 - Today: Are we focusing our optimization efforts on the right pages? 7
 - Looking ahead: New questions 7
 - 1. How can we better measure how performance affects user satisfaction? 7
 - 2. What impact does web performance have on customer lifetime value (CLV)? 7
 - 3. What impact does performance have on enterprise productivity? 8
 - 4. Are we always measuring the right things? 8
 - Six takeaways 8
- The Performance Beacon 10
 - Downtime vs slowtime: Which hurts more? 10
 - Visitors are more likely to permanently abandon a slow site than an unavailable one 11
 - Slow pages could have twice the impact on revenue that site failures do 12
 - Slow sites suffer more damage to brand health 12
 - Takeaway: Preventing outages is just one piece of the performance pie 12
- Conversion Impact Score: What is it? And why do you need to know yours? 14
 - First: What does “conversion” mean? 14
 - Performance slowdowns affect conversions differently on different pages 16
 - What is the Conversion Impact Score? 17
 - Case study: How to use Conversion Impact Scoring to prioritize performance optimization 17
 - Mistake #1: Prioritizing the Order Billing page because it’s the slowest 19
 - Mistake #2: Tackling the Shopping Bag and Home pages next 19
 - Mistake #3: Not worrying about the Product and Category pages because they seem relatively fast 19
- Conclusion 19
- Takeaway 19

A brief history of web performance ROI

A brief history of web performance ROI



It was apparent to anyone who was paying attention that performance really does touch pretty much every metric the business cares about.

2010: Should we care about page slowdowns as much as we care about downtime?

Outages get a lot of attention, for obvious reasons. When twenty minutes of downtime can mean millions in lost revenue... well, that's pretty scary. Because of this, it's sometimes been a challenge to get business owners to care about slow load times as much as they care about downtime.

Fortunately, the folks at TRAC Research came to the rescue with [an excellent study into the real impact of outages versus slowdowns](#). They surveyed more than 300 companies and found that, on average:

- 4.4 seconds is the average delay in response times when business begins to be affected.
- \$21,000 is the average revenue loss for one hour of downtime.
- \$4,100 is the average revenue loss for an hour of slowdowns.
- Website slowdowns occur 10 times more frequently than website outages.

In other words, **website slowdowns have twice as much negative impact on revenue as outages do.**

2011: Can we compare apples to apples?

Many performance case studies are makeover stories. We see a “before” version of a slow site, then an “after” of a revved-up version of the same site. The “after” version almost invariably results in happier users and better business metrics.

But these “before and after” stories inevitably have detractors who make some good points. Most sites aren’t static. Perhaps the uptick in business metrics is due to a marketing campaign? Or perhaps some of the credit should go to changes in the user interface? **Before and after stories are compelling (after all, most of us love a web performance Cinderella story), but they don’t convince everyone. That’s why we need to compare apples to apples, instead of apples to oranges.**

In other words, we need to see the impact of performance changes on pages where all other elements are the same and the only changing variable is load time. But doing this kind of research is difficult.

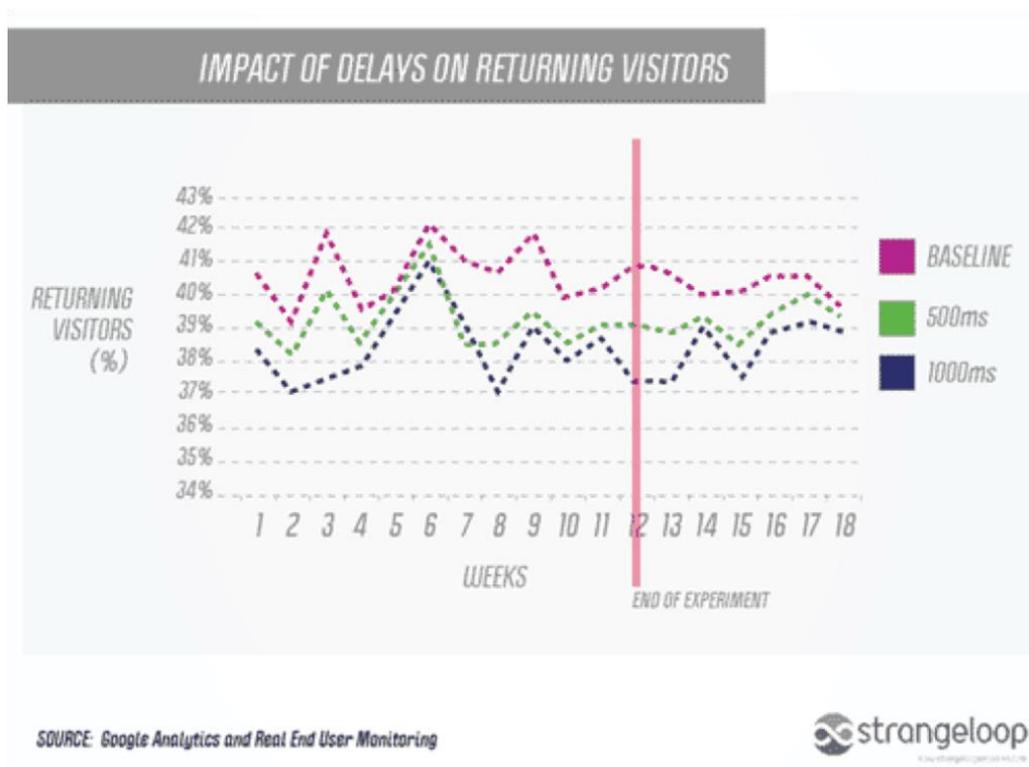
I feel very lucky to have participated in [a somewhat unusual study](#) back when I was with Strangeloop. We had a customer who wanted to see firsthand the impact of page slowdowns on their own site, so they agreed to an 18-week study in which, for the first 12 weeks, the bulk of their traffic was served an optimized version of the site, but three small (but still statistically significant) cohorts were served pages with either a 200 millisecond, 500 millisecond, or 1000 millisecond HTML delay.

The results were eye-opening. While the 200-millisecond delay had a negligible impact on business metrics, at 500 milliseconds the impact was quite pronounced. **And at 1000 milliseconds, we saw an 8.3% increase in bounce rate, a 3.5% hit to conversion rate, a 2.1% decrease in cart size, and a 9.4% drop in page views.**

	BOUNCE RATE	CONVERSION RATE	CART SIZE	PAGE VIEWS
200 ms	—	—	—	-1.2%
500 ms	-4.7%	-1.9%	—	-5.7%
1000 ms	-8.3%	-3.5%	-2.1%	-9.4%

— NO SIGNIFICANT CHANGE

Not only were business metrics affected during the 12 weeks of the study in which we introduced the HTML delay, they continued to be affected even after we resumed optimizing 100% of site traffic. We continued to monitor the visitors from the first 12 weeks of our study for 6 more weeks, and we found that the percentage of visitors who returned was lower for the cohorts that had received the 500 millisecond and 1000 millisecond delays. **In other words, the slower load times these users experienced negatively affected their willingness to return to the site, even after the experiment was over.**



2012: What can we do with a LOT of user data?

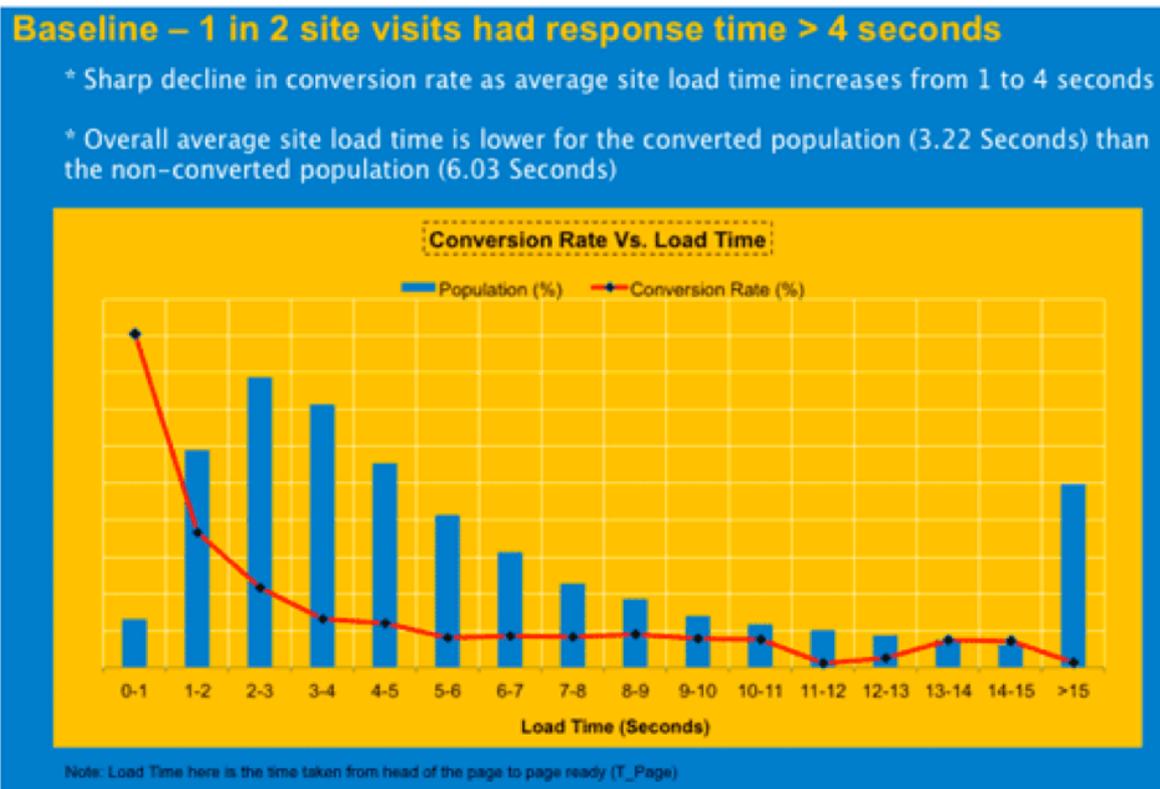
2012 was notable for being the year that real user monitoring (RUM) came on the scene. While people were generally excited by the idea of gathering data about every user experience on their site, we also wondered how the heck to take all this data and make it actionable.

One of the earliest RUM case studies that showed up on my radar was from Walmart Labs, a tech incubator within Walmart. (The research was led by Cliff Crocker, who served as senior engineering manager at Walmart Labs.

The team at Walmart Labs knew that Walmart.com was suffering from performance issues. As a for instance, initial measurement showed that an item page took about 24 seconds to load for the slowest 5% of users. Why? The usual culprits: too many page elements, slow third-party scripts, multiple hosts (25% of page content was served by third parties), and various best practice no-nos.

Walmart Labs dedicated a scrum team to one sprint of performance optimization. At the start of the process, the team performed some baseline measurements in which they used their real user monitoring data to look at the load times of key pages and look for patterns. Then the team created targets for page performance and, at the end of the sprint, measured the impact of optimization on key metrics.

Impact of site performance on overall site conversion rate....



Their findings showed a strong correlation between performance and conversions:

- Overall, converted shoppers were served pages that loaded twice as quickly as pages served to non-converted shoppers.
- This trend persisted, even on individual pages that experienced greater load times.
- Non-buyers were served category pages that were 2-3 seconds slower than category pages served to buyers.
- **For every 1 second of improvement to load time, the site experienced up to a 2% improvement in conversion rate.**
- For every 100 milliseconds of improvement, they grew incremental revenue by up to 1%.

Today: Are we focusing our optimization efforts on the right pages?

“We worked to make our home page and other key pages faster, and it didn’t affect our business.”

I don’t hear this very often, but I do hear it. One reason why some companies work hard to optimize key pages and yet still not see results: they’re optimizing the wrong pages. Some pages (such as product and category pages) are more sensitive to performance changes than others (such as checkout pages).

Knowing this, our data science team here at SOASTA developed something called the Conversion Impact Score, which answers the question: What’s the relative impact of load time changes on business performance per page? Knowing the Conversion Impact Scores for pages on your own site lets you prioritize your optimization efforts. It’s a powerful metric. Every time I talk about it with our customers, I can see the excitement it generates.

Looking ahead: New questions

When I look back and see how much research has been generated and how many questions have been answered over the past seven years, it’s really exciting. But I’m even more excited about the questions we haven’t answered. Here are just a few things I’m looking to explore:

1. How can we better measure how performance affects user satisfaction?

We have tools that measure performance. We have tools that measure customer satisfaction. It stands to reason that we can have these tools speak to each other so that we can visualize the relationship between site speed and user happiness.

2. What impact does web performance have on customer lifetime value (CLV)?

In the performance monitoring space, we tend to measure the user experience in single sessions. But that’s a really short-sighted way of looking at things. In the marketing world, marketers take a big-picture look at the entire customer relationship. Typically, CLV is calculated over three years: that’s the length of the relationship that marketers expect a customer to have with their brand. So when marketing calculates the ROI for bringing in a new customer, they look at the amount of revenue that customer brings in over the entire relationship.

We in the web performance community need to align our metrics around customer experience with marketing's metrics. Rather than focusing solely on single user sessions, it would be incredibly helpful to gather performance data over a much longer window of time.

3. What impact does performance have on enterprise productivity?

Much of the research on the business impact of web performance focuses on retail, but when it comes to using enterprise apps, performance affects workers, too. Thanks to apps like Slack, workers have high expectations when it comes to speed and ease of use. We need to gain a better understanding of how app performance affects internal metrics like app adoption and employee satisfaction and productivity.

4. Are we always measuring the right things?

As Steve Souders said at WebPerfDays Amsterdam:



Tammy Everts @tameverts · Oct 31

"Most of us are not measuring the user experience correctly."

@Souders #WebPerfDays #webperf #perfmatters

2:31 AM - 31 Oct 2015 · Details

At Velocity Amsterdam, there was a lot of talk about using UserTiming as a metric for gathering much more specific data around how pages render and how we measure user satisfaction and other business metrics. This discussion bled into WebPerfDays, with some deep discussion into how to use UserTiming in the field. As Steve also said:



Tammy Everts @tameverts · Oct 31

"We're going to see companies figure out how to automatically collect user timings from their RUM data." @Souders #WebPerfDays

#webperf

RETWEETS

6

LIKES

4



2:59 AM - 31 Oct 2015 · Details

>>

Six takeaways

When I talked about performance ROI in my webinar, my hope was that people would leave having internalized these messages:

1. User expectations and behavior are always changing
2. Our tools for measuring ROI are evolving

3. Know your own business success metrics
4. Understand your own visitors
5. Target the best (not the slowest) pages for optimization
6. Monitor, test, repeat

The Performance Beacon

The Performance Beacon

Downtime vs slowtime: Which hurts more?



Tuesday's AWS outage – which lasted four hours and affected almost 150,000 companies, including a significant number of online retailers – is, not surprisingly, being dissected to bits.

Downtime is horrifying for any company that uses the internet as a vital part of its business (which is to say: most companies). It was almost exactly one year ago, in March 2016, that Amazon (the retailer, not the service) famously went down for 20 minutes. [Those 20 minutes may have cost the company \\$3.75 million in lost sales.](#)

That's a big number taken by itself – no company wants to think about losing millions in revenue – but it's a drop in the bucket for a company with a [net revenue of more than \\$125 billion](#) in 2016. While Amazon assuredly takes pains to avoid outages, the

company also goes to great effort to manage the day-to-day performance – in terms of page load – of its retail site. That’s because Amazon knows that page slowdowns can cause at least as much harm as downtime.

There are three metrics that are hit harder by slow page loads:

- Abandonment rate
- Revenue
- Brand health

Let’s take a deeper dive into the data behind each of these metrics.

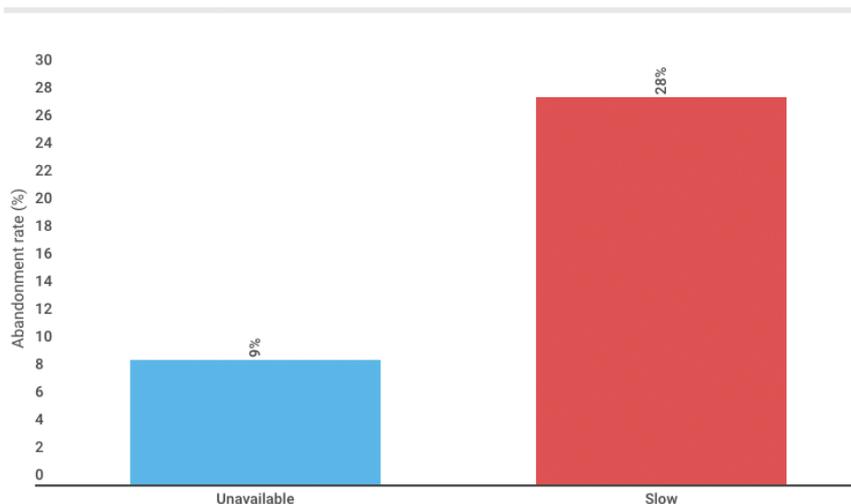
Visitors are more likely to permanently abandon a slow site than an unavailable one

If a website is temporarily down, there’s a reasonable chance you’ll try again later if you’re motivated to track down whatever it was you were interested in finding on that site. But if a website or app is consistently laggy (read: many popular media sites), eventually you just sort of drift away.

Anecdotally, this makes sense, and there’s research to back it up.

In one of the only studies (if not *the* only study) of the impact of outages versus slowdowns on abandonment rates, Akamai found that sites that went down experienced, on average, a permanent abandonment rate of 9 percent. **Sites that suffered from slow performance experienced a 28% abandonment rate** – an increase of more than 200 percent.

Permanent abandonment rate



Source: Akamai, The Impact of Web Performance on E-Retail Success

This isn’t to say that site outages are nothing to be concerned about. A 9% permanent abandonment rate is extremely detrimental to your business. And a 28% abandonment rate is even worse.

Slow pages could have twice the impact on revenue that site failures do

This finding comes from a study that, to the best of my knowledge, is the only study that compares revenue losses due to downtime with losses due to page slowness. TRAC Research surveyed 300 companies and found that the average revenue loss for an hour of downtime was \$21,000. For the same set of companies, average revenue loss due to an hour of performance slowdown (which was defined as response times exceeding 4.4 seconds) was much less (just \$4,100).

Looking at just these two sets of numbers, outages seem like a bigger source of concern. But wait. According to this same survey, **website slowdowns occurred ten times more frequently than outages**. In other words, according to this research, slow-loading pages could have twice the impact on revenue that site failures do.

Slow sites suffer more damage to brand health

Unless your site experiences frequent and noticeable outages, occasional failures won't undermine your brand. Most users accept sporadic downtime as part of the reality of using the web. They're less forgiving, however, if your site is routinely slow.

First impressions matter, and they happen faster than you might think. According to one study, we form our opinion of a website within the first 50 *milliseconds*. And once we've formed that opinion, it colors how we feel about a site's credibility and usability, ultimately affecting whether or not we choose to make a purchase on that site.

A few years ago, I directed a neuroscientific research project in which participants were asked to complete transactions on an ecommerce site using mobile devices. Some participants experienced normal speeds, while others experienced load times that were artificially throttled with a 500-millisecond network delay. Participants believed they were participating in a standard usability/brand perception study, so they had no idea that speed was a factor in the tests.

After each set of tests, researchers conducted exit interviews with the subjects, who were asked to give their general impressions of each site and company. The results were revealing. Some participants picked up on the slight deterioration in performance (calling the slower site "slow" and "sluggish"), but those who used the slower site also developed negative perceptions of areas unrelated to speed:

- Content ("boring")
- Visual design ("tacky" and "confusing")
- Ease of navigation ("frustrating" and "hard to navigate")

In other words, **the slower pages affected people's perception of three important aspects of the site that are closely aligned with brand perception.**

Takeaway: Preventing outages is just one piece of the performance pie

There's no such thing as 100% uptime. Every site goes down eventually. The question isn't "*Will my site go down?*" The questions are "*When will my next outage happen?*" and "*How long will it last?*" and "*How much will it cost me?*"

If your business is reliant on your site, then you most definitely should care about preventing outages. You should, of course, conduct load testing and performance testing, and you should have effective load balancing and failover systems in place wherever possible.

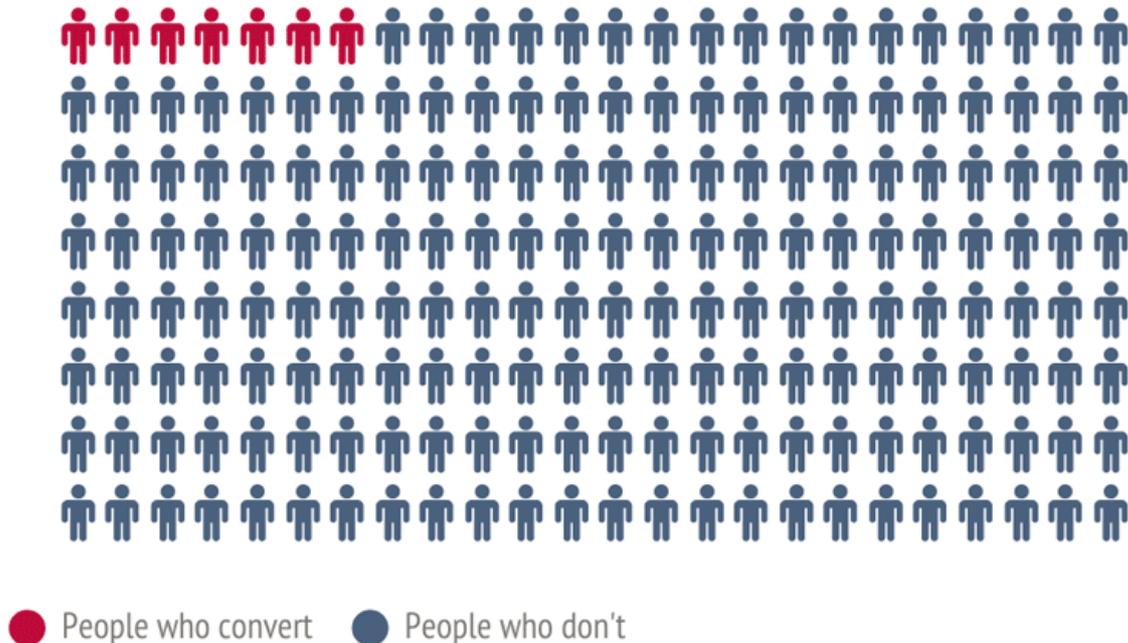
But protecting your site from failure is just one piece of the performance pie. It's a big piece, to be sure, but there are others. You also need to:

- **measure** your site's performance;
- **correlate** IT performance metrics (start render, load time) with business and user engagement metrics (conversions, bounce rate, session length) so that you know which IT metrics are optimal for your business;
- **monitor** performance in realtime to ensure you're hitting your optimal IT metrics;
- **drill down** and find and resolve performance issues as they occur; and
- **look for opportunities to further optimize** your pages (hint: images and third parties are a great place to start).

Comparing outages to slowdowns is like comparing a tire blowout to a slow leak. One is big and dramatic and the other is quiet and insidious. Either way, you end up stranded on the side of the road.

Conversion Impact Score

Conversion Impact Score: What is it? And why do you need to know yours?



From a user experience perspective, in an ideal world every page would load in less than a second, we'd zip through a transaction in moments, and boom, we'd be done.

But as I've written about in the past, not all web pages are created equal. People react differently to slowdowns on different pages in the conversion funnel, which means you need to approach each page differently. While it would be wonderful if we could optimize every single page of our websites, most site owners have only a finite amount of optimization resources. You need to focus those resources on optimizing the pages that matter most to your bottom line.

In this post, I'm going to explain how to determine which pages you should focus on optimizing in order to increase conversions and, ultimately, deliver the fastest ROI.

First: What does "conversion" mean?

There's a widely held belief that the only people who need to care about conversions are people in sales and business development. Wrong. Conversions are the lifeblood of your business. If you touch your company's website in any way — be it design, marketing, or development — then your actions have an impact on conversions. You need to understand what that impact is.

But first, I have a tendency to drop terms like “conversion” and “conversion rate” without explaining what they mean, so let me slow down for a minute to offer some definitions.

Conversion = What happens when a person who’s browsing a site converts to being a user or buyer of the service or product that site offers. So if you’re a SaaS vendor, a conversion happens when a person signs up to use your service. Or if you’re an ecommerce shop, when a person buys something. Conversions can also include actions like signing up for a newsletter or making a donation. You get the idea.

Conversion funnel = The start-to-finish path that a person takes when they convert from browsing to buying/downloading/etc. A conversion funnel for an ecommerce site might look something like this (note that percentages are arbitrary):



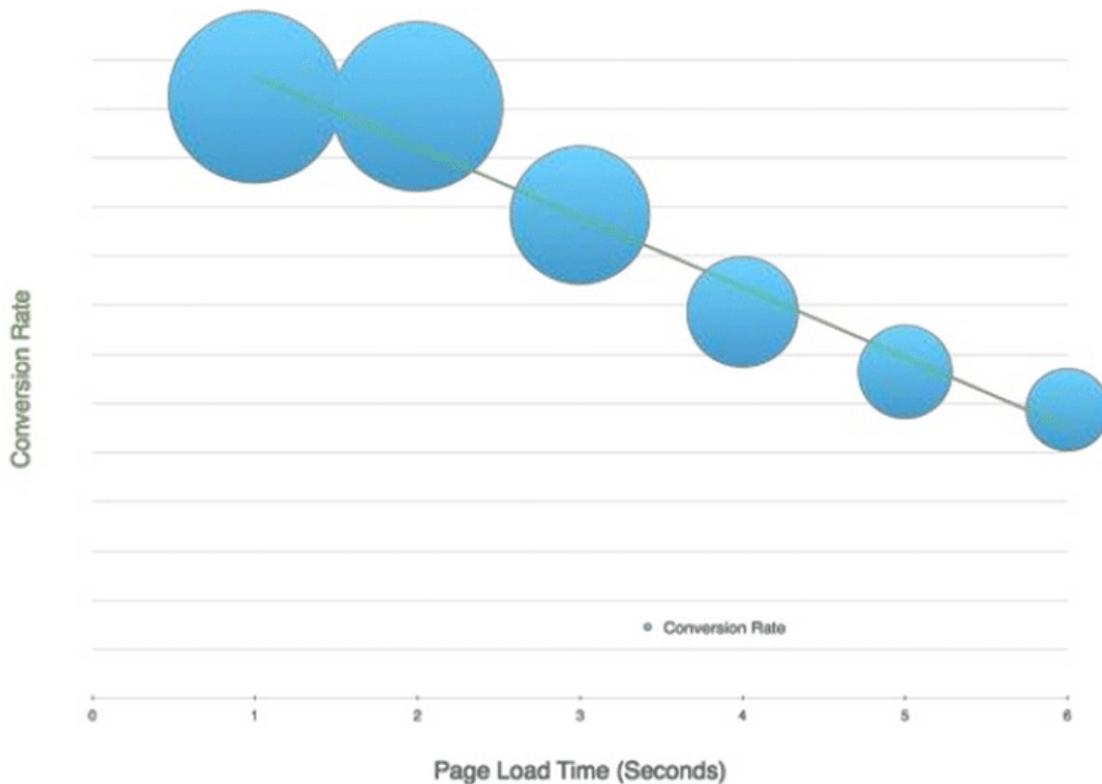
Conversion rate: The percentage of total visits to a site that result in a conversion. In the conversion funnel graphic above, the conversion rate is the number of people who completed a purchase: 10%. Note that 10% would be a wildly successful conversion rate. More typically, conversion rates are in the 2-4% range. Anything higher than that is amazing. For a site that does hundreds of thousands of dollars worth of transactions in a day, even tiny changes in conversion rate — such as increasing from 2.1% to 2.2% — can have a huge impact on revenue.

Performance slowdowns affect conversions differently on different pages

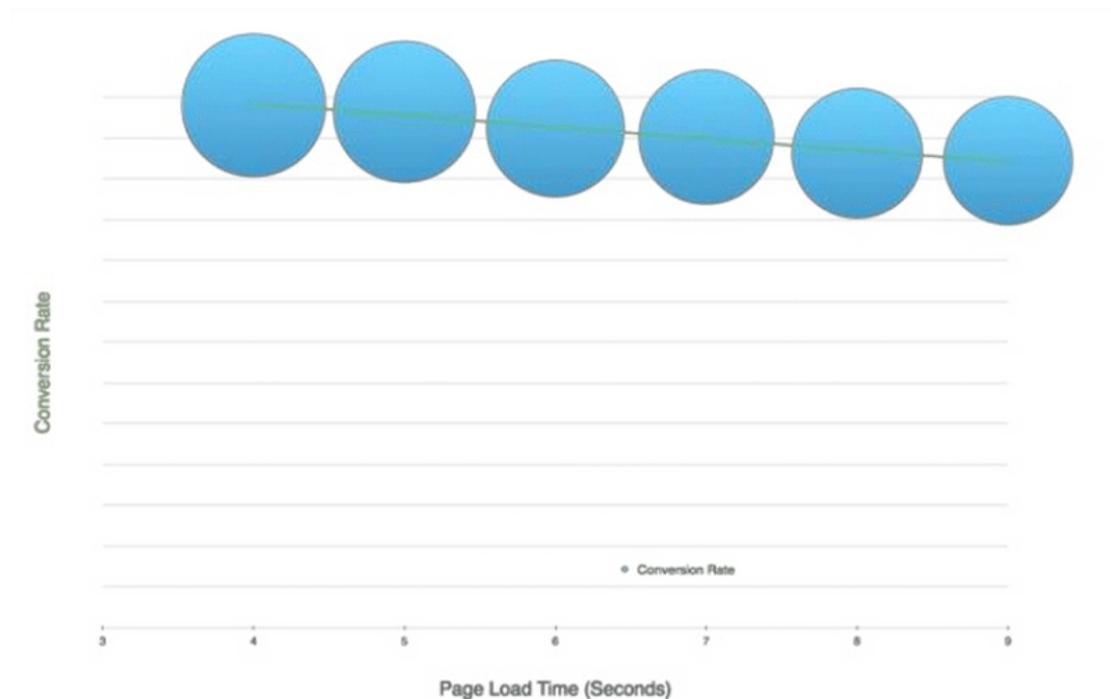
When pages get slower, conversion rates suffer. But some types of pages suffer more than others.

A while back, we looked at real-life performance data for one of our ecommerce customers who uses real-user monitoring. We found that, while slower load times correlated to fewer conversions, the impact was most dramatic when pages in the “browsing” part of the conversion funnel were slower.

Here you can see that the conversion rate shrinks by about 50% when the load time for “browse” pages increases from 1 to 6 seconds:



Looking at the same set of user data, you can see that the impact on conversion rate is much less when checkout pages degrade in speed:



Looking at these two graphs side by side, you could be tempted to deduce that, because conversions were hurt more by slow “browse” pages than by slow “checkout” pages, the site owner should focus energy on optimizing the browse pages. This might be true — but it might not be true, too. This is where the Conversion Impact Score comes in.

What is the Conversion Impact Score?

The long definition:

The Conversion Impact Score (CIS) is a relative score that ranks page groups by their propensity to negatively impact conversions due to high load times. For each page group, the Conversion Impact Score is calculated using the proportion of overall requests that are associated with that group, along with the [Spearman Ranked Correlation](#) between its load times and number of conversions. The Conversion Impact Score will always be a number between -1 and 1, though scores much greater than zero should be very rare. The more negative the score, the more detrimental to conversions that high load times for that page group are, relative to the other page groups.

The TL;DR definition:

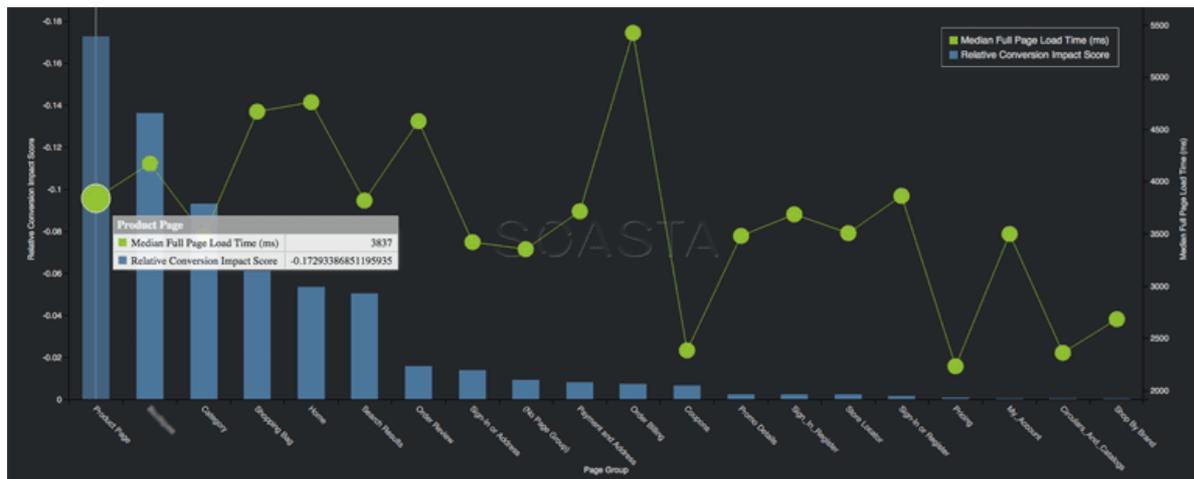
The Conversion Impact Score answers this question: How much impact does the performance of this page have on conversions?

Case study: How to use Conversion Impact Scoring to prioritize performance optimization

Now let's walk through how to use the Conversion Impact Score to make decisions about optimizing your pages.

In the graph below (which represents anonymized data from one of our mPulse customers), you can see the Conversion Impact Scores and load times for a set of pages on a site. The blue bars represent the Conversion Impact Score for each page, and the green line represents the median page load time for each page.

The pages are ranked from those with the highest Conversion Impact Scores (such as product and category pages; in other words, pages viewed in the "browse" phase of the conversion funnel) to pages with the lowest scores (such as sign-in and account pages).



Some quick observations:

- The three fastest pages — Coupons, Pricing, and Circulars and Catalogs — all have relatively low Conversion Impact Scores, despite being very fast. This means that page speed isn't a significant factor in how well these pages convert.
- The pages with the best Conversion Impact Scores — such as Product and Category pages — have okay load times in the 4-second range.
- The slowest page is the Order Billing page (in other words, a checkout page), followed by the Home page, Shopping Bag, and Order Review pages.

Without knowing the Conversion Impact Scores for these pages, you might focus on optimizing pages according to how slow they load. Looking at load time, this is the order in which you'd prioritize fixing these pages:

1. Order Billing
2. Home
3. Shopping Bag
4. Product
5. Category

Now here's how these assumptions are incorrect:

Mistake #1: Prioritizing the Order Billing page because it's the slowest

If you looked only at page load times, you might believe that you need to prioritize the Order Billing page because its performance is dramatically poorer than the other pages. But if you knew its Conversion Impact Score, you'd realize that page speed doesn't have much impact on conversion rate either way, so making this page faster wouldn't be the best use of your limited optimization resources.

Mistake #2: Tackling the Shopping Bag and Home pages next

Also, if you were to look exclusively at load times, you might believe that when you're done with optimizing the Order Billing page, you should focus next on addressing performance issues on the Shopping Bag and Home pages. While these pages have a high enough Conversion Impact Score that they merit addressing, they shouldn't rank high up on your list.

Mistake #3: Not worrying about the Product and Category pages because they seem relatively fast

Still looking solely at load times, you might also guess that, because the Product and Category pages look fairly speedy, you don't need to worry about them. This is where you'd make your biggest mistake. Because these pages have the highest Conversion Impact Scores, they have the potential to deliver the most benefit to you if you make them faster.

Conclusion

So, knowing the Conversion Impact Scores for this set of pages, this is the order in which you might actually want to prioritize their optimization to give you the best ROI:

1. Product
2. Category
3. Shopping Bag
4. Home
5. Order Billing (This is debatable. While improving the performance of this very slow page might not have much impact on conversions due to its low Conversion Impact Score, it would definitely improve the user experience at the tail end of the transaction — which is when people are most impressionable about their perception of an experience. Improving the final stage of the transaction could result in a higher degree of customer satisfaction, which could improve customer retention and word of mouth. This is all somewhat theoretical, but it has its basis in other research around performance and user experience.)

Takeaway

Knowing your pages' load times is just one piece of the optimization puzzle. You need to correlate load time with other metrics, such as conversions or bounce rate, that are meaningful to your business, and you need to develop a tool — such as the

Conversion Impact Score — that can weight your results so that you know what pages to tackle first to get the best ROI for your optimization efforts.