17 PERFORMANCE METRICS YOU SHOULD CARE ABOUT

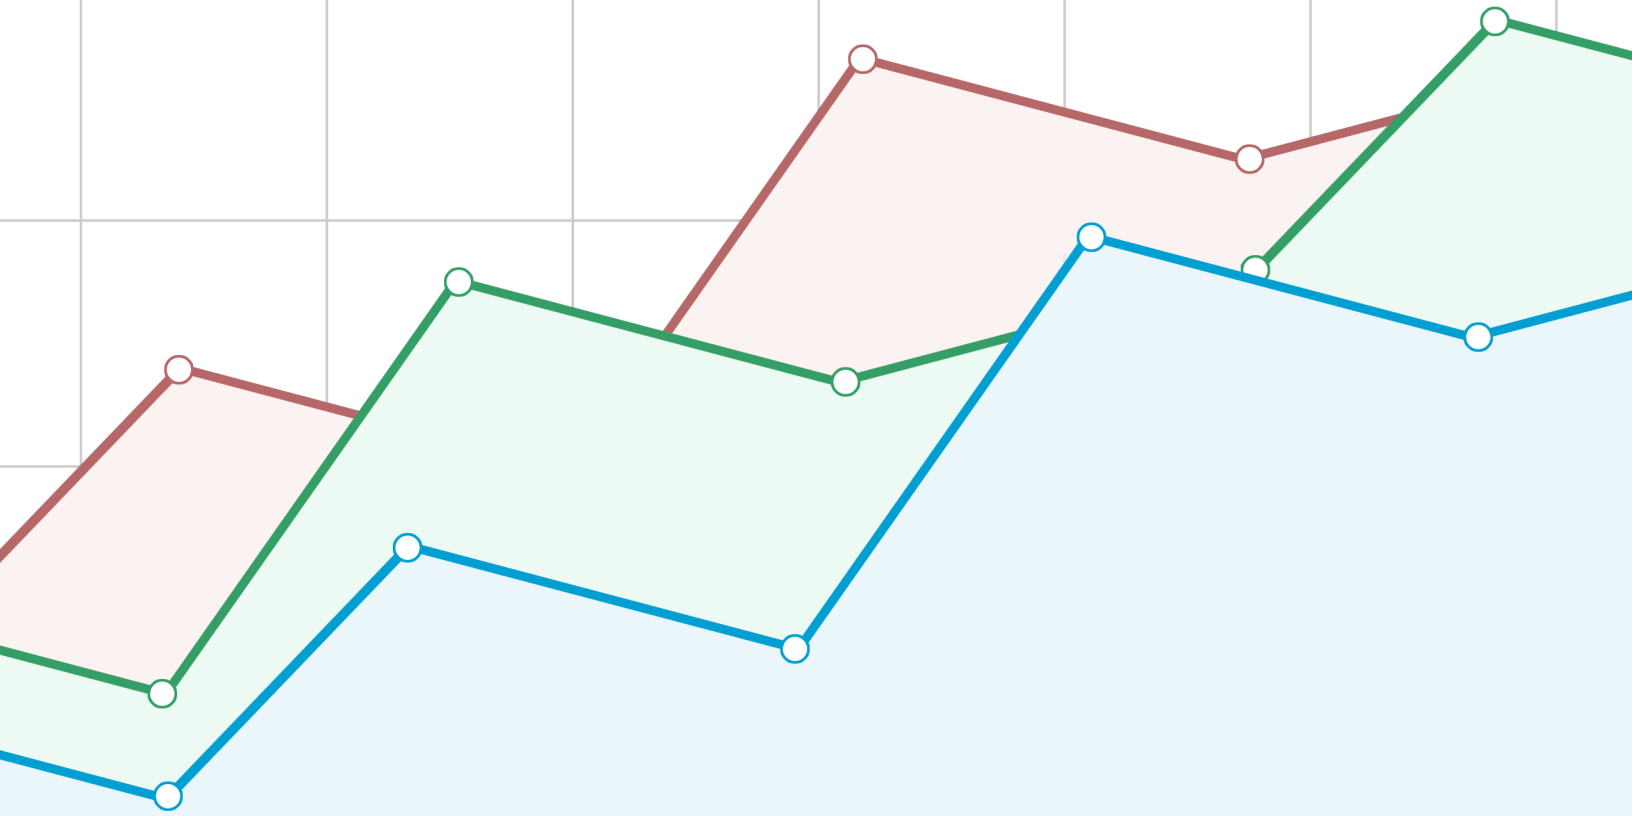

yottaa

contents

“if you can’t measure it, you can’t improve it”

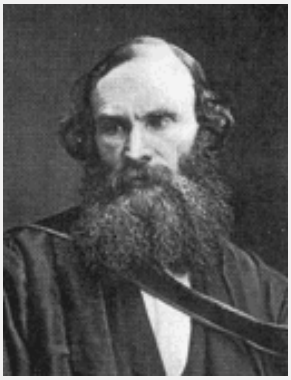

(Lord K. thrives on performance metrics)

Creating a great user experience (UX) is among the most challenging aspects of running a website today. On one hand, users today expect websites to offer tons of information, media, and interactivity – all with page load times of two seconds or less. On the other hand, the Internet’s growing complexity makes it ever-harder to achieve great performance. There are simply more ways to go wrong than ever before (see below).

At Yottaa, we see site owners every day struggle to strike a balance of great performance and content-rich sites. Our message to these site owners is almost always the same. We tell them that the best tool for dealing with the UX challenge is – wait for the cliché – *knowledge*.

In the world of web performance optimization, knowledge is gained through performance monitoring. Site owners looking to improve performance should use monitoring to attain intimate familiarity with every part of their website. Collecting performance data and interpreting the results properly will unlock the door to truly great website user experience. As our good friend and famed 19th century scientist Lord Kelvin says, “If you can’t measure it, you can’t improve it.”

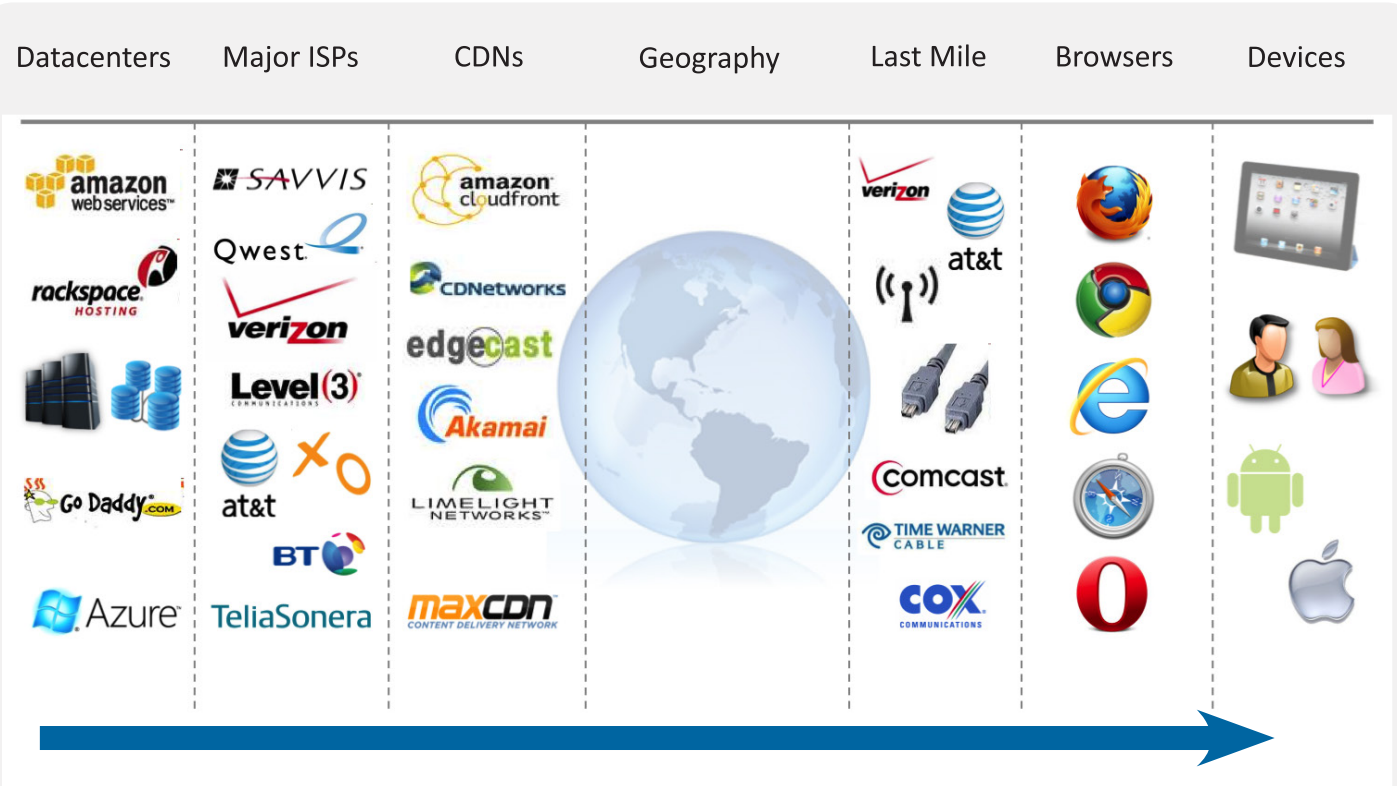

Delivering a quality experience for all users gets more challenging with every additional variable. As you can see, web delivery today involves a number of variables at each step of the way.

In this eBook we will cover 17 metrics that we think any site owner will find useful. We'll describe why each metric is important, and share average figures for each from a survey of thousands of websites monitored by the Yottaa network. We hope reading this eBook will give you the foundation necessary to set up a performance monitoring service for your site, understand the results in a meaningful way, and benchmark your site against the global averages, as well as best-in-class stats within your target market.

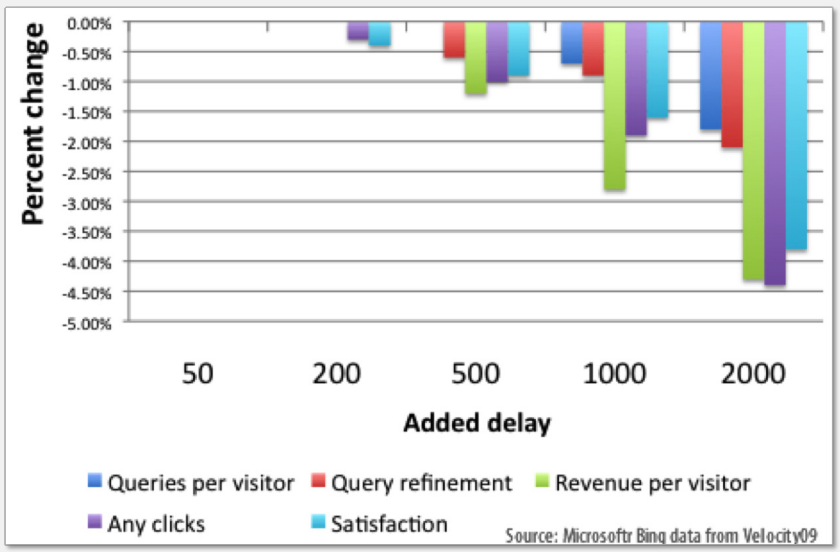

This graph from Microsoft's Bing team shows the decline in key user metrics as page speed delay is added

Keep this in mind: when we talk about "monitoring" we mean more than getting a single line of data across these metrics. In order to achieve statistically significant data you will need to test multiple samples from multiple locations around the world. The Internet is dynamic beast, with dozens of factors aligning to deliver your site and a plethora of devices and browsers in play on the client-side. Knowing whether or not your site is consistent for all visitors across these many variables is key to successful web performance optimization.

Juggling these multiple variables requires "multi-variate testing." A certain amount of testing is required before setting up monitoring to ensure that the story told by the continuous stream of monitoring data is true. Once you have established that your site performs with consistency across different browsers and different last-mile connectivity speeds, you can monitor with confidence. Familiarity with the 17 metrics we present here will equip you not only to understand your monitoring data, but to do this preliminary testing as well.

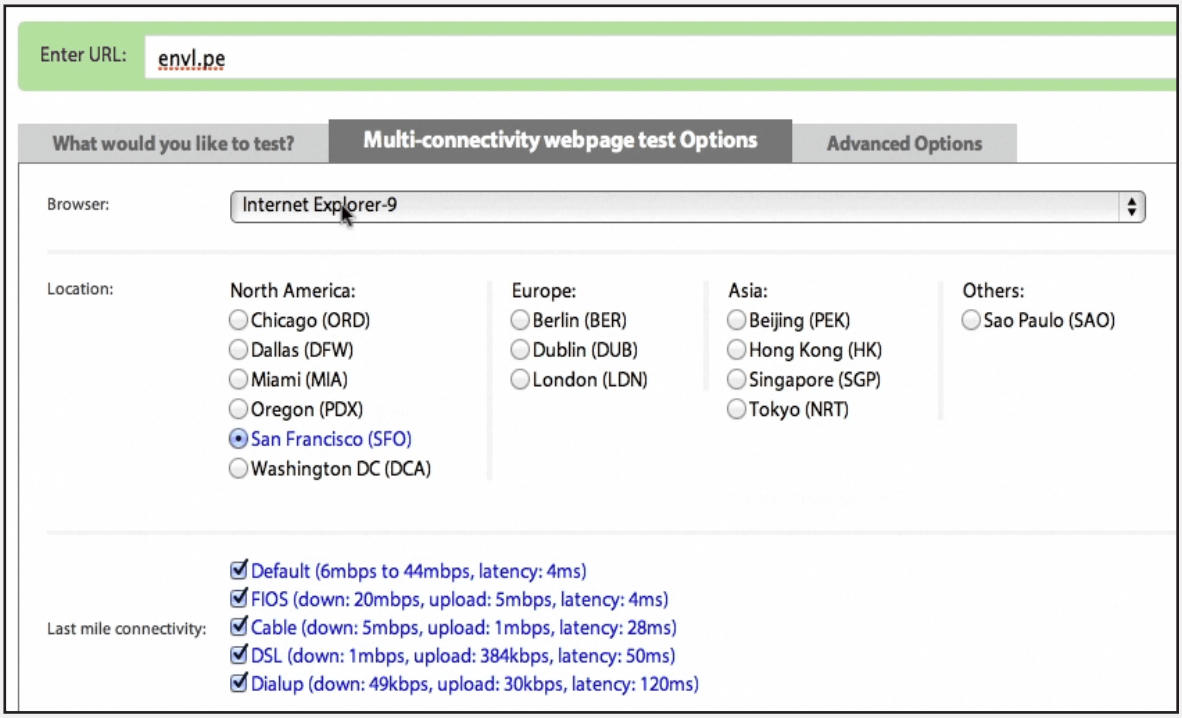

overview: the basics

Broadly, the 17 metrics covered in this eBook fall into three categories: front-end user experience, backend performance, and content complexity.

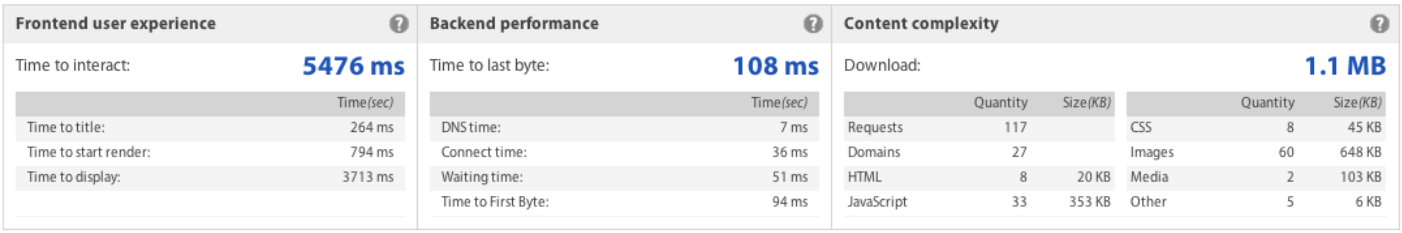

Frontend user experience

Time to interact: **5476 ms**

	Time(sec)
Time to title:	264 ms
Time to start render:	794 ms
Time to display:	3713 ms

Backend performance

Time to last byte: **108 ms**

	Time(sec)
DNS time:	7 ms
Connect time:	36 ms
Waiting time:	51 ms
Time to First Byte:	94 ms

Content complexity

Download: **1.1 MB**

	Quantity	Size(KB)
Requests	117	
Domains	27	
HTML	8	20 KB
JavaScript	33	353 KB

	Quantity	Size(KB)
CSS	8	45 KB
Images	60	648 KB
Media	2	103 KB
Other	5	6 KB

Front end user experience refers to how the user's browser executes your page's content. Aspects of this process are apparent to the visitor, such as when your site's title appears in the browser tab, when your site's first visuals are displayed in the frame, and when the visitor can interact with your site. These are important metrics to understand because they have the most impact on a user's perception of page load time, which is often just as important as the load time itself!

Backend performance refers to how your site gets to the visitor's browser in the first place. Everything that happens from the moment a visitor clicks the link to your site to the moment the last byte of your site arrives at her browser is encompassed in this category. Your hosting provider, DNS provider, and content delivery network (if you have one) all factor into these metrics.

Content complexity is a breakdown of the building blocks of your website. It's important to know what media and programming languages are present on your site, how many assets there are of each, and how heavy those assets are. These are not performance metrics per se, but the content of your site has an immediate impact on front-end metrics, so it's important to understand this aspect.

front-end metrics

Four stages in the front-end execution process are commonly noted: Time to Title, Time to Start Render, Time To Display, and Time to Interact. (These stages are known by a few different names, but for simplicity we'll stick to these.) Each of these is noticeable to the visitor, so changes in these metrics will not only affect your nominal page load time, but will also directly affect user experience. When bounce rates go up or down while your content hasn't changed, look to these metrics for potential answers and solutions.

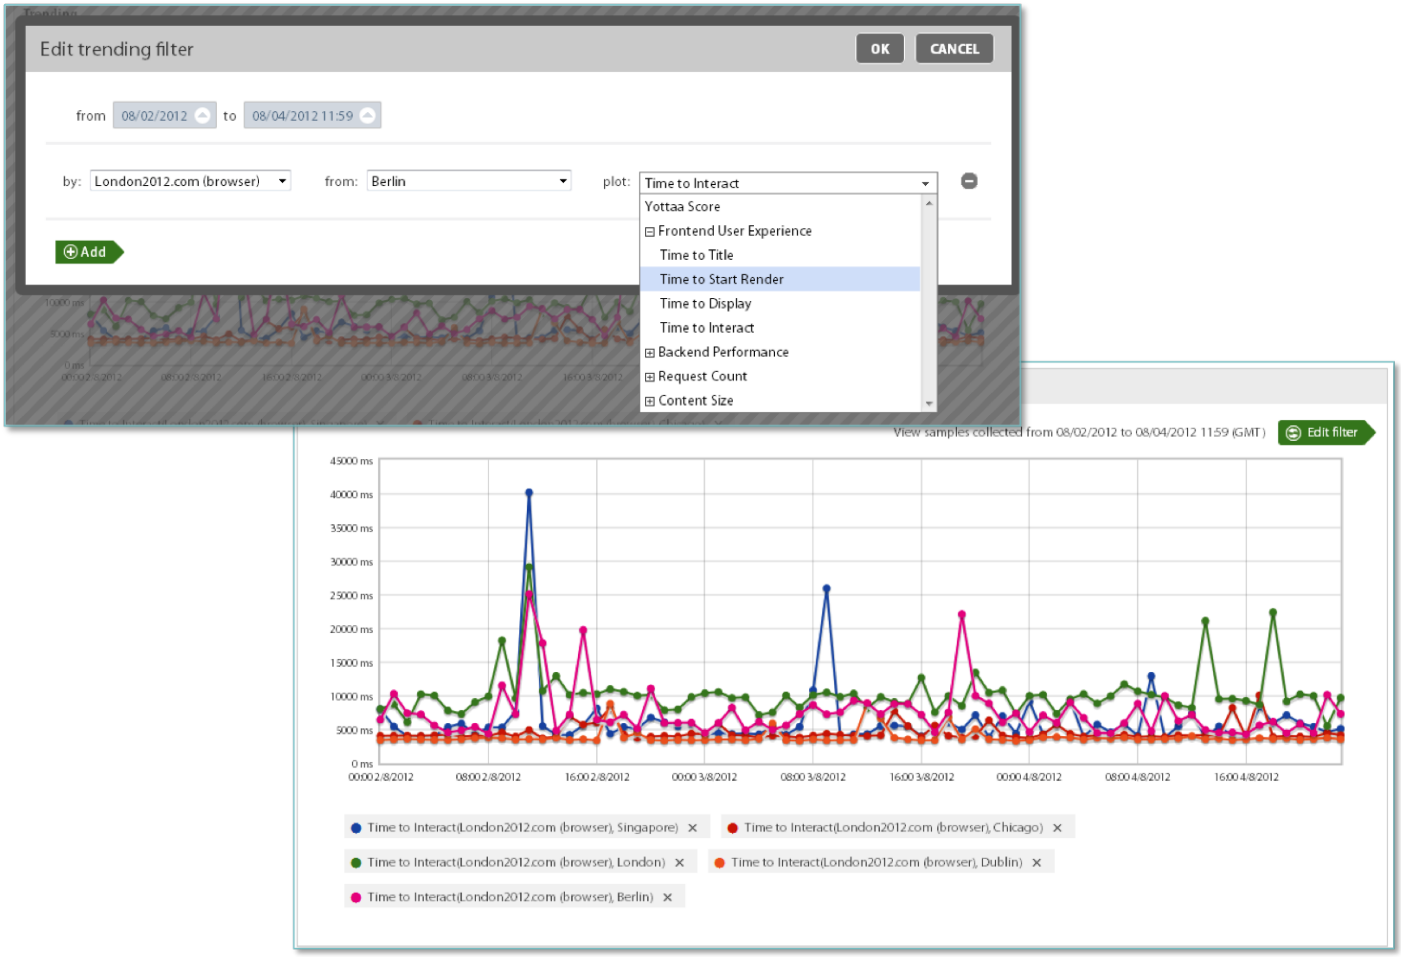

These images show how you can view monitoring statistics in Yottaa Monitor. The trending graph displays data filtered to your specifications.

Metric 1

Time To Title

Time To Title is the time elapsed from the moment a visitor's browser requests your site to the moment that your site's title appears in the tab of his or her browser.

If you're already familiar with this metric then you know that it actually has little to do with the front end of your website. Time To Title is largely dictated by your backend performance—that is, how fast your website's content is delivered from your origin server to the visitor's browser. So why is it in the front-end user experience category rather than backend? Because on the front-end, where user perception of load time is half the battle, Time To Title is the first indication to the visitor that your site exists and is on the way. When users see the title appear promptly they are apt to be patient throughout the load process. This makes Time To Title the first of several user experience metrics.

	50th percentile	80th percentile	95th percentile
Time to Title (msec)	1,259	2,368	4,647

Note - 50th percentile refers to the median value for this metric (and all that follow). That is, half of the measurement sample (>2,000 sites) had a Time to Title faster than 1.26 seconds, and half the sites were slower. If your site's Time to Title is 2.37 seconds or slower, it's among the slowest 20%; and if it's 4.65 seconds or slower, it's in the slowest 5%.

Metric 2

Time To Start Render

Time to Start Render is the time elapsed from the request to when the visitor sees actual website content appear on the page. Like the Time To Title, this moment is important since it assures the visitor that your site is loading. Nobody likes staring at a blank page. Assuring visitors that they are in the right place and will be soon seeing the content they expect will promote a good perception of your website.

	50th percentile	80th percentile	95th percentile
Time to Render (msec)	2,420	4,071	7,160

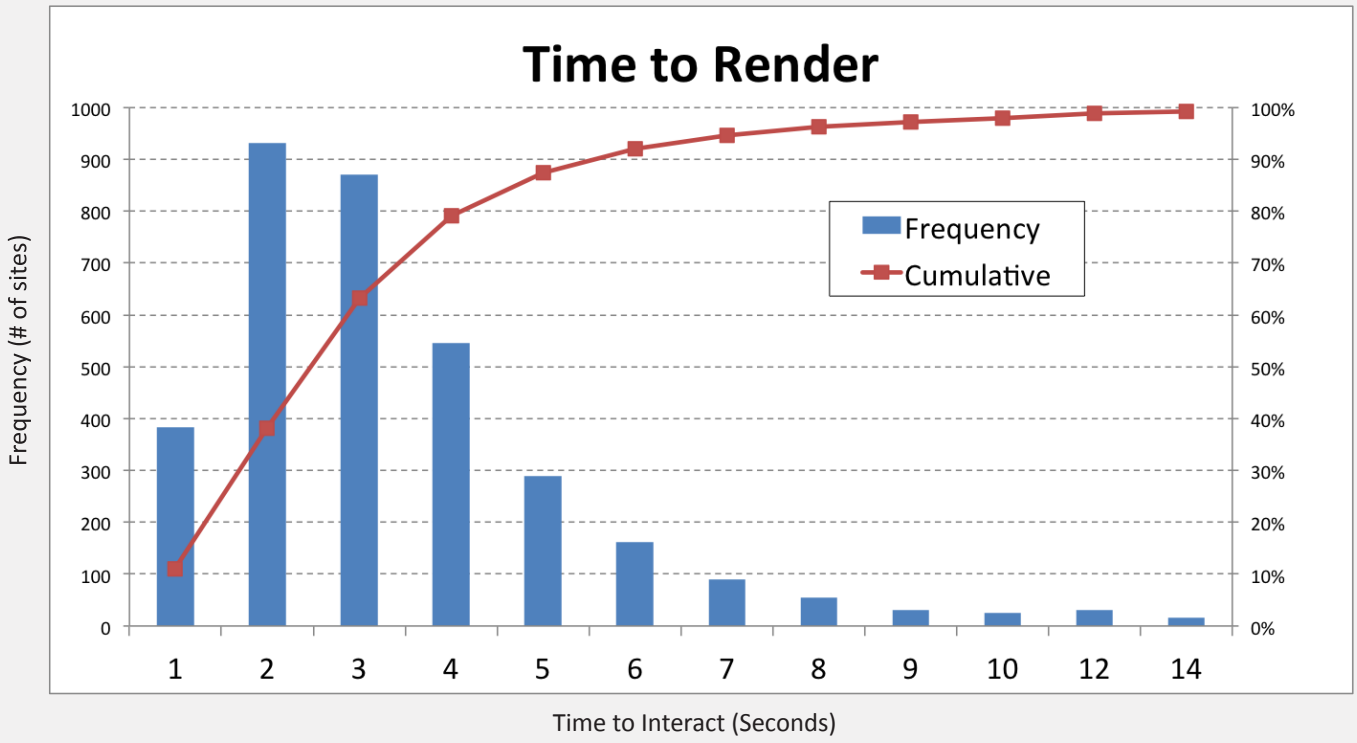

Metric 3

Time To Display

Time To Display is the time elapsed from the request to when the browser has finished parsing the HTML page, constructed the Document Object Model (DOM), and displayed the HTML document. This all means that the page will look like a web page, but there may be some images, interactive elements, and other media that haven't fully loaded.

	50th percentile	80th percentile	95th percentile
Time to Display (msec)	5,168	8,780	15,254

Metric 4

Time To Interact

Time To Interact is the time elapsed from the request to the moment the user can interact with the page. (By "interact" we mean the page will respond properly to the visitor clicking a link, scrolling, typing into a field, or activating an element like a hover effect). This does **not** mean that the page is fully loaded, as there may be scripts, trackers, and other assets that continue to load in the background. But it does mean that the almighty visitor can use the web page, and that's an important moment. Many site owners choose Time To Interact as the principle index for overall web performance because of its relationship with user experience.

	50th percentile	80th percentile	95th percentile
Time to Interact (msec)	6,263	10,643	18,494

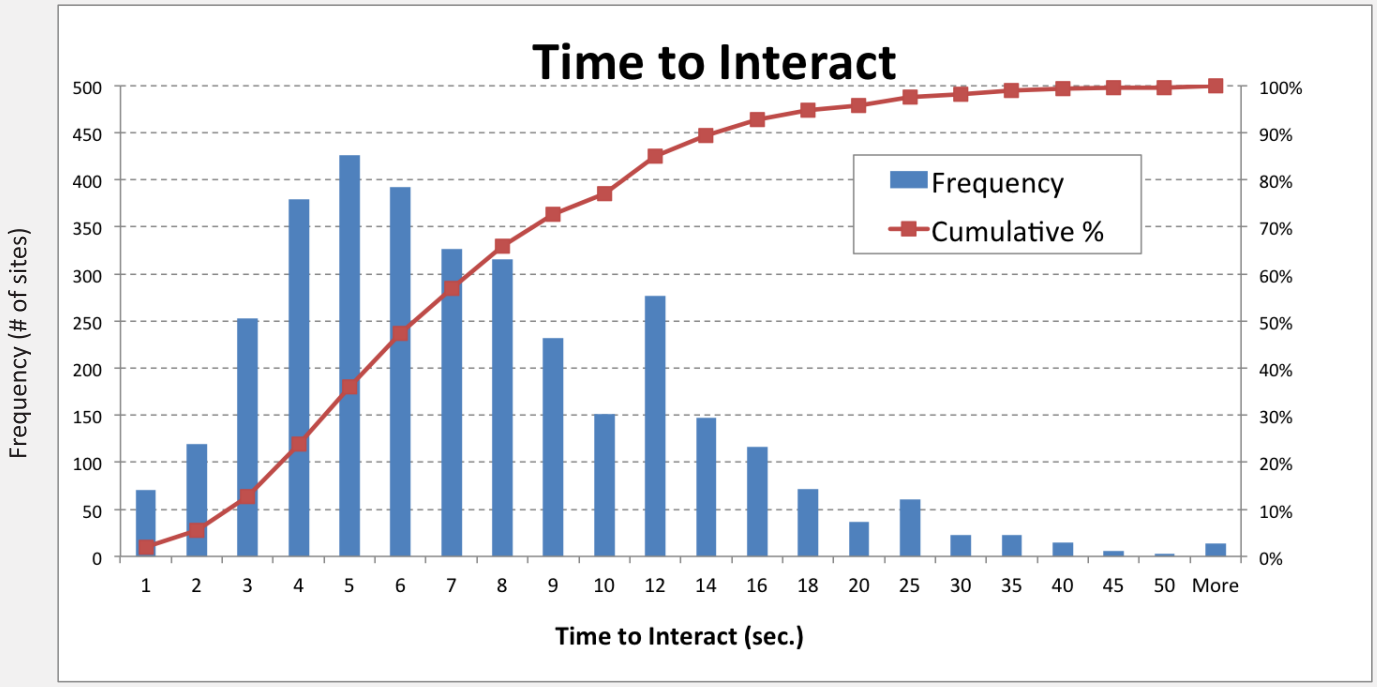

backend metrics

Backend processes can be broken down into stages, much like the front end. Time in the backend is more compressed, however: the last stage in backend delivery typically takes less than a second to complete. Plus the stages in backend delivery are not directly visible to the visitor, so there's less emphasis on optimizing each stage from a user experience standpoint. That doesn't mean that the backend can be neglected, however -- if backend performance is slow, the entire front-end process will be delayed!

If you're looking at a waterfall chart (a chart that graphs the load time of each individual asset of a website, found on webpagetest.org and websitetest.com) the backend performance can be seen visually in the first line. This is also known as the HTTP Timeline. On websitetest.com, mousing over this timline will bring up the details on the following metrics (see below).

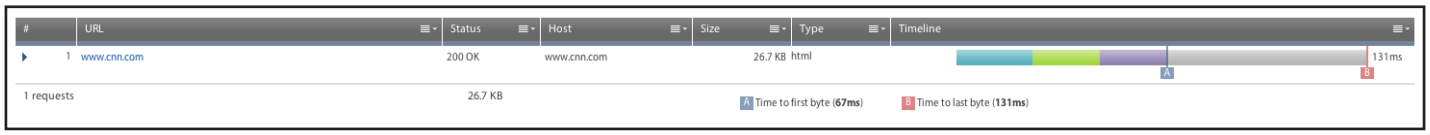

An HTTP timeline for CNN.com via Websitetest.com

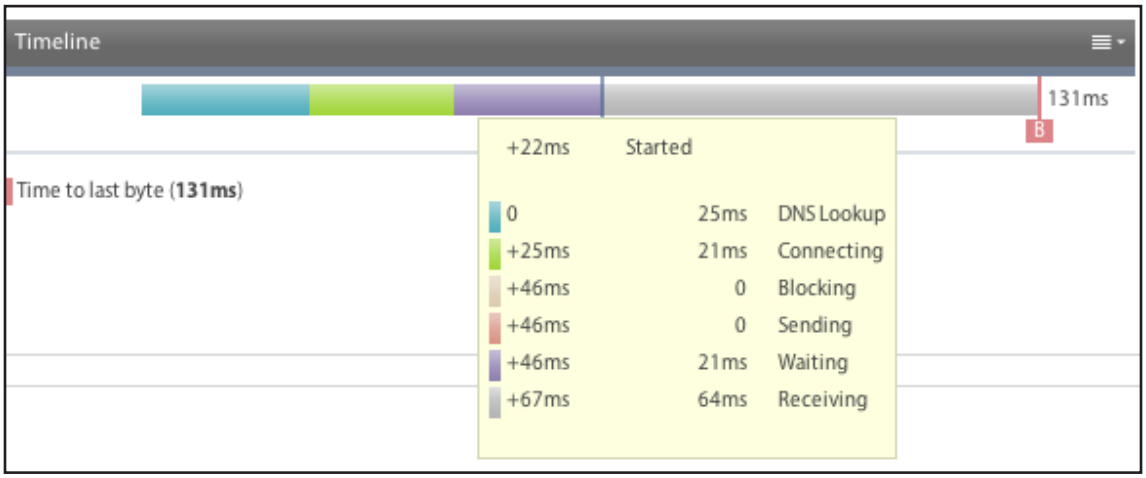

The backend metrics shown when mousing over the timeline for CNN.com via Websitetest.com

Metric 5

DNS Time

DNS Time is the time it takes for your DNS provider to execute its service. (A DNS processes the visitor's URL request and returns the IP address that matches the URL by searching through enormous databases). You can find your DNS time with an online service like DNSStuff.com or WebsiteTest.com.

A typical DNS time is 60 milliseconds. If your DNS time is much slower than that, you may be best off finding a new DNS provider. A good rule of thumb is to find an IP anycast-based Global DNS service (Yottaa provides one for free!).

	50th percentile	80th percentile	95th percentile
DNS Time (msec)	60	149	413

Metric 6

Connection Time

Connection time is the time elapsed from the request to when the connection between the visitor's browser and your origin server is established. Normal connection times are around 150 milliseconds.

Problems with connection time can be a bit more difficult to diagnose because they only crop up in certain circumstances. If your server is overloaded with traffic -- either because of heavy traffic on your site or because of bots – connection times will spike. Geography can also play a role, with distant visitors experiencing longer connection time.

Since these problems may not happen frequently, continuous monitoring may not be enough to spot the problems. Load testing with tools like LoadStorm or JMeter will simulate increased traffic and reveal the increased connection time and/or 503 and 504 errors that would occur in times of heavy server usage.

If load testing reveals serious issues, the solutions are to upgrade your infrastructure or offload some of your assets onto a content delivery network (CDN) or caching server.

	50th percentile	80th percentile	95th percentile
Connect Time (msec)	144	188	227

Metric 7

First Byte Time

After the connection is established, the next step in the backend process is for the first byte of your website to travel along the connection to the visitor's browser. Once the first byte arrives, your First Byte Time is registered.

The only impediment to achieving a good First Byte Time is poorly structured code. Nearly every web page has content that is "static" and does not need to be personalized to the particular visitor. This content should be separated from the rest so that it can be sent from the server to the visitor immediately, rather than waiting for personalized content to be processed and sent along with it. Issues with First Byte Time can be found with a combination of continuous performance monitoring and load testing.

	50th percentile	80th percentile	95th percentile
Time to First Byte (msec)	533	1,090	2,418

Metric 8

Last Byte Time

Last Byte Time is the final backend metric. It is logged when every byte of your website has made it to the visitor's browser. An average last byte time is about 700 milliseconds, although we've seen times well over three seconds in the data for the higher percentiles.

Like First Byte Time, the factors in determining Last Byte Time are largely to do with the code and database. If your code is not optimized, your database queries are poorly written, or if you're not taking advantage of memcache, your Last Byte Time could be suffering.

As with the rest of the backend metrics, finding Last Byte Time is easy: just pick a monitoring service. However, digging into the reasons behind your Last Byte Time may require more effort. Load testing tools, as well as application performance management (APM) tools like Gomez, Keynote, and Yottaa, may be necessary to complete the picture.

	50th percentile	80th percentile	95th percentile
Time to Last Byte (msec)	721	1,412	3,105

content complexity

In order to put front end and backend metrics to use, you'll need to know about the composition of your website. Content complexity statistics can reveal areas of your site with potential for improvement. Just as with timing metrics, you can use percentiles from our sample of the Web to determine how the composition of your site stacks up against the rest.

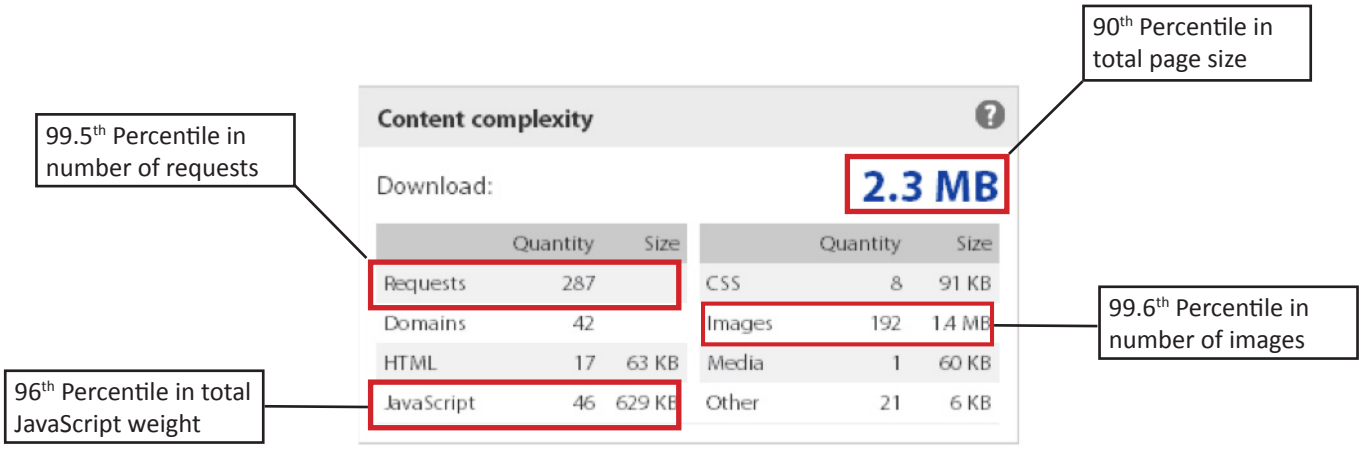

Above is content complexy data from NBCOlympics.com captured during the London 2012 summer games on WebSiteTest.com. The site is slow -- in the 81st percentile of our web survey with a Time To Interact over 10 seconds -- and the content complexity helps tell us why. Notice the extremely high percentiles in four major content categories -- no wonder the site is so slow!

Metric 9

Total Asset Weight

Metric 10

JavaScript Weight

Metric 11

CSS Weight

Metric 12

Image Weight

The overall weight -- that is, the number of bytes -- of your website factors into its speed, as does the weight of individual assets (one heavy and slow-loading asset can have ripple effects on performance).

Use weight metrics to identify categories of assets that are too heavy in aggregate, and then use a waterfall chart to zero in on specific assets within that category that can be fixed or cut.

For instance, if the total weight of your JavaScript files is in a much higher percentile than your other content categories, click to the waterfall chart in your monitoring service. If you're in Yottaa Montior, you can click Type > JavaScript to show only JavaScript files, then Size > Descending to show the heaviest JavaScript files at the top.

	50th percentile	80th percentile	95th percentile
Asset Weight (Bytes)	690,779	1,557,812	3,393,192
JS Weight (Bytes)	119,378	346,852	586,442
CSS Weight (Bytes)	18,134	60,265	120,000
Image Weight (Bytes)	318,788	895,948	2,056,237

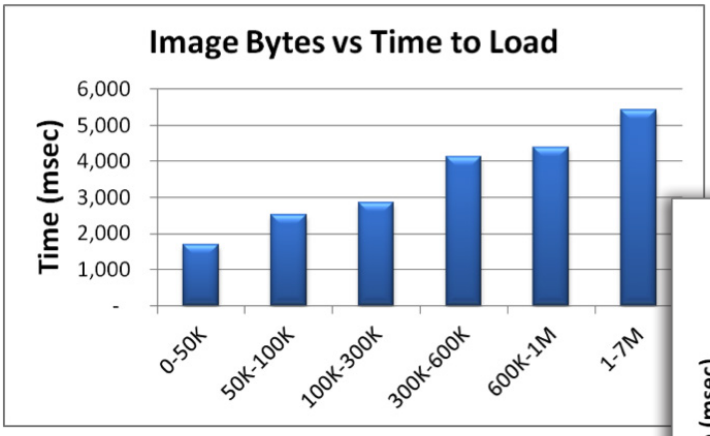

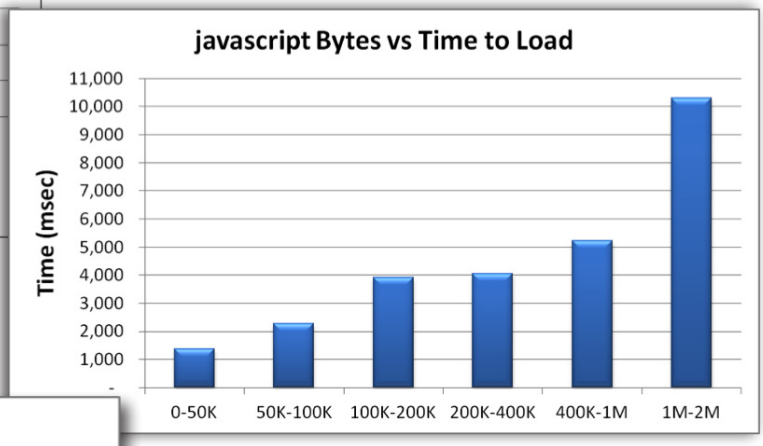

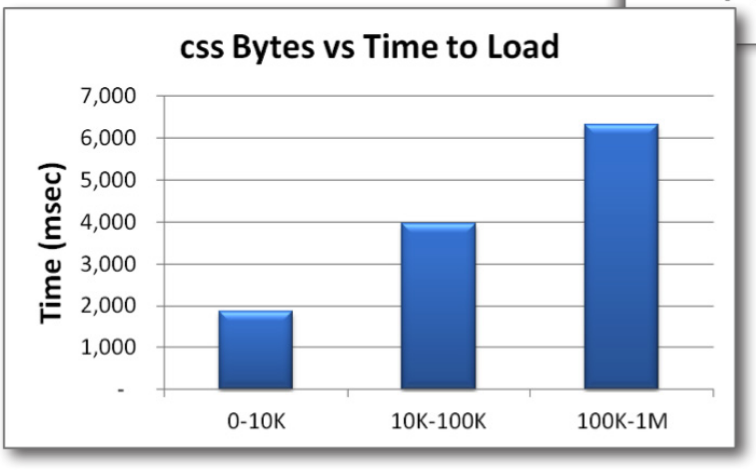

Metric 13

Overall Asset Count

Metric 14

JavaScript Count

Metric 15

CSS Count

Metric 16

Image Count

More assets necessarily mean more weight -- that's reason enough to keep track of asset count. But in addition, each time a visitor's browser makes the trip to your origin server to fetch an asset for your site, it adds time to the page load. That means each asset slows down your site no matter how small or compressed it is.

	50th percentile	80th percentile	95th percentile
Asset Count	47	84	152
JS Count	8	19	33
CSS Count	3	8	16
Image Count	25	48	87

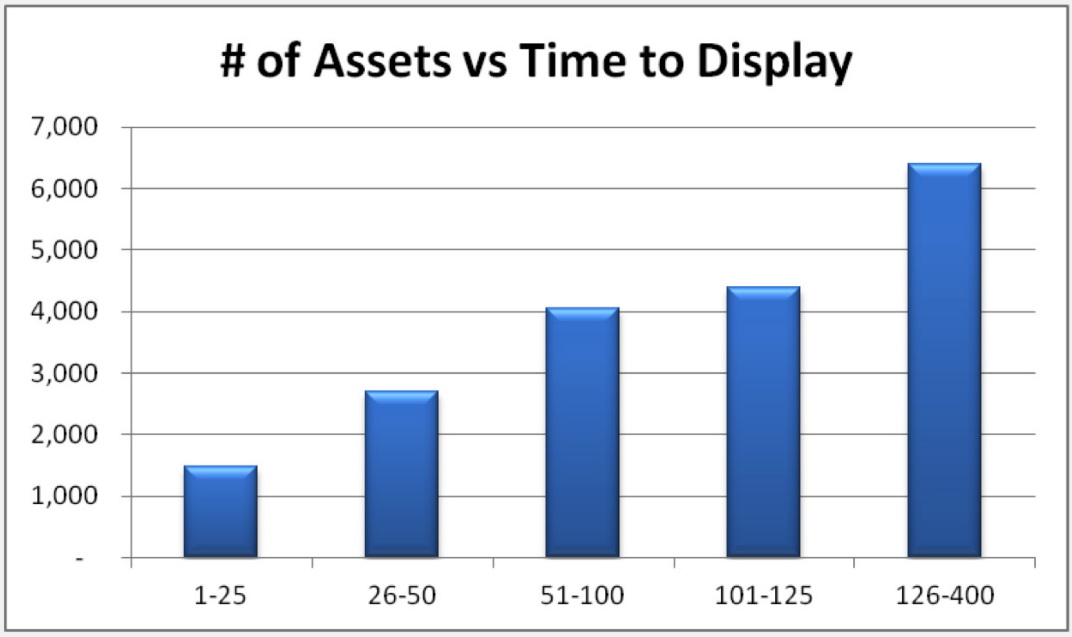

Metric 17

Domains

Most of your site's content is hosted on your domain (or "origin") server. But if your site features content from a third party, such as widgets for social media interaction or an embedded YouTube video, then there are additional domains in the mix. Taking content from other domains would not be bad for performance in a perfect world -- but in reality, the inconsistent nature of the web means more domains means more risk for performance issues. If one of your third party content providers has an issue, it could drag down your entire site. Monitor your domains to be sure your site is not relying too heavily on multiple other sources.

	50th percentile	80th percentile	95th percentile
Domains	6	17	33

the next level - monitoring tips

There's more to monitoring than metrics. Some services offer additional functionality that augment and automate the monitoring process. These are not necessary to achieving knowledge about your website's performance, but they save you time and offer peace of mind -- so why not put them to use?

Tip

Establish and Track SLAs

Wherever there's monitoring data, there's opportunity for a service level agreement (SLA). Based on the results of your prelimiary performance testing, set up objectives for your website's uptime and performance. These SLAs will be tracked along with your continuous monitors, letting you keep track of big picture goals as you navigate your stream of monitoring data.

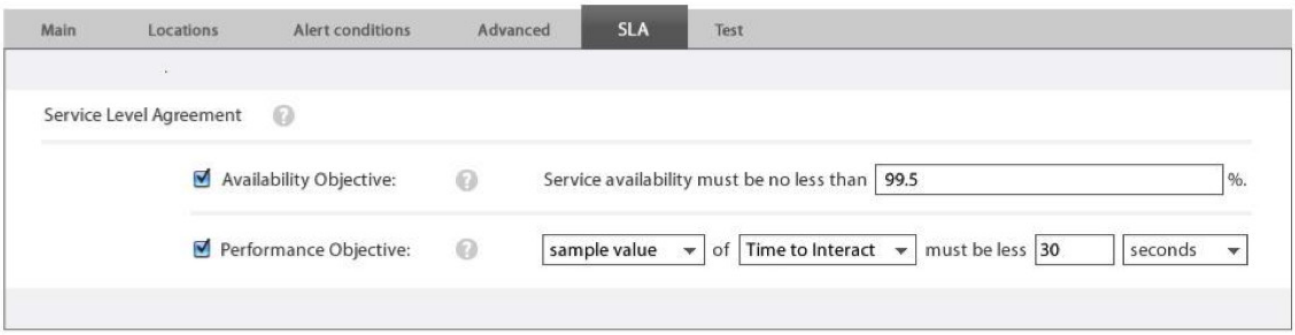

Tip

Set Up Issue Definitions

As your monitoring service collects data, it will inevitably detect problems on your pages. Setting up issue definitions lets you decide how the monitoring service presents these issues to you -- whether a given issue is presented as a critical error, as a warning, or not noted at all. (With Yottaa monitor you are set up with some basic issue definitions by default, but these can be altered or deleted.)

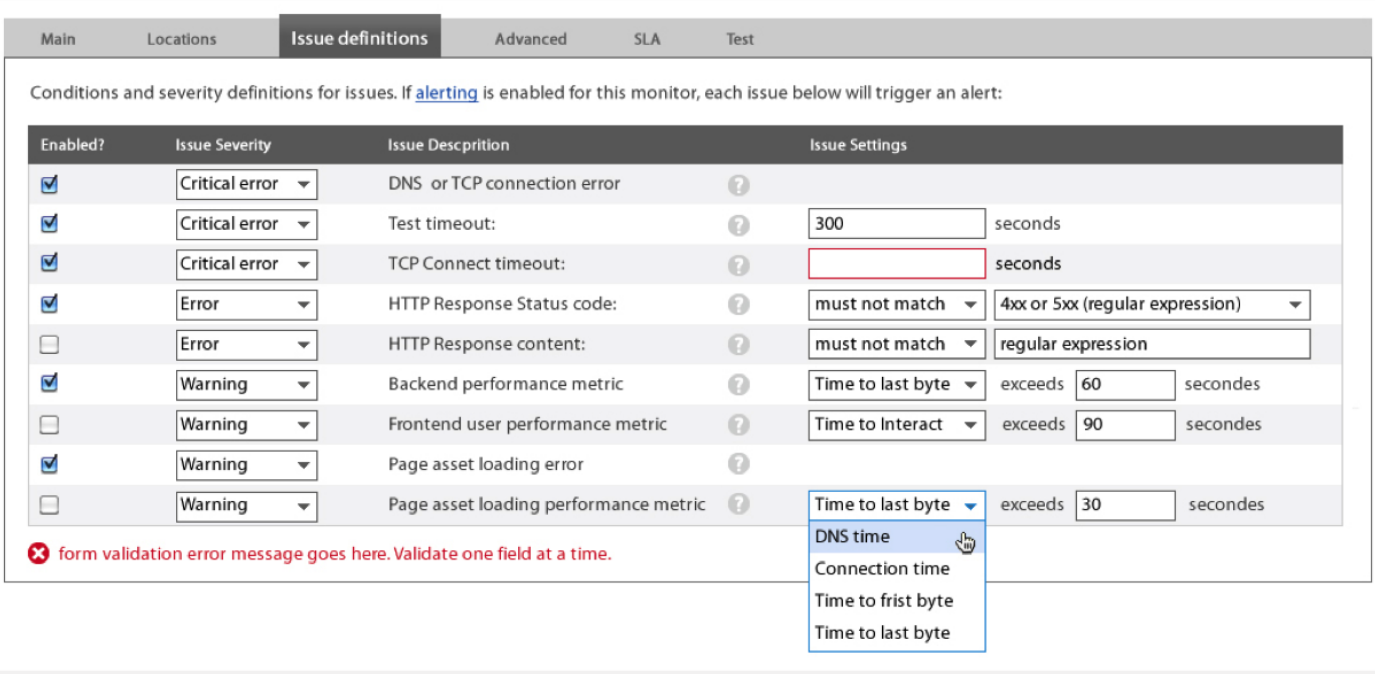

Tip

Set Up Alerts

Alerts are the logical followup to setting issue definitions. You can arrange for emails, or in some cases text messages, to be sent to you and your coworkers when the issues you defined come up.

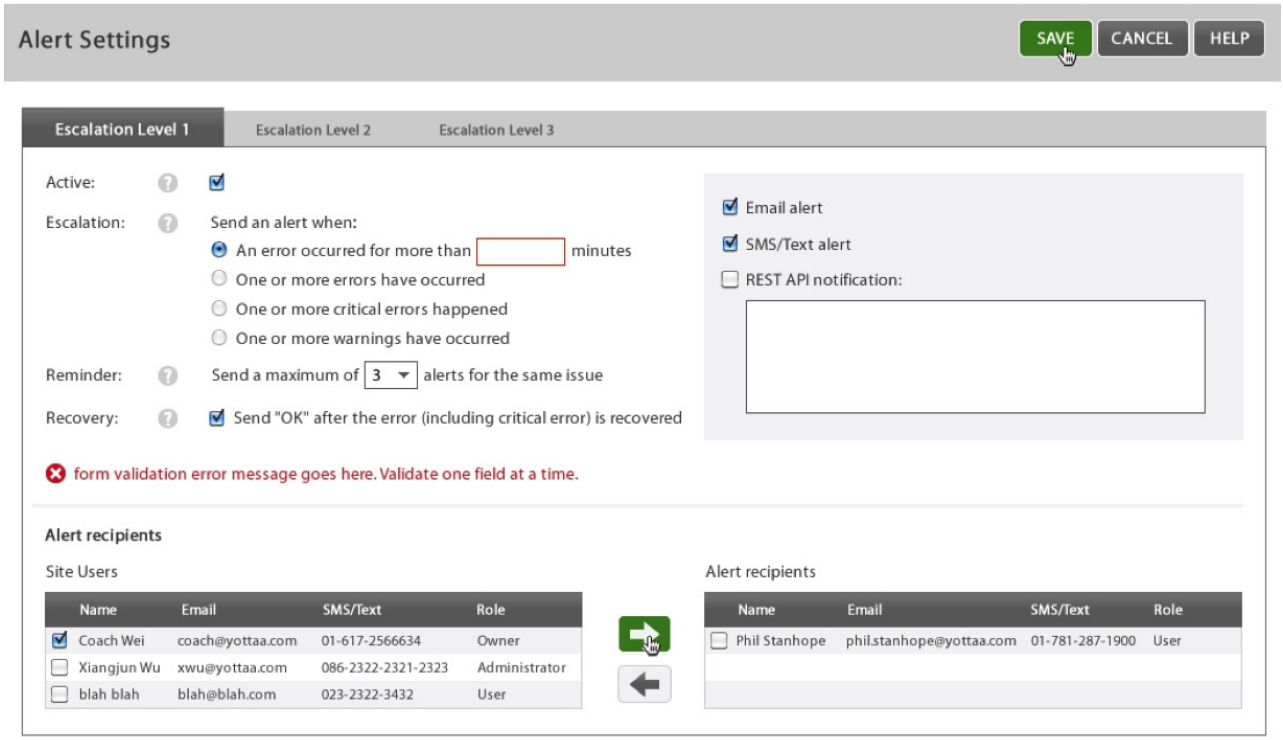

conclusion

Now that you have the baseline information under your belt, it's time to get to work. Find out how your site performs on all the metrics described in this eBook with a free Yottaa Site Monitor account. There's no time limit on the trial and it will get you started on the path to effective monitoring!

Try it Free

Yottaa Site Monitor

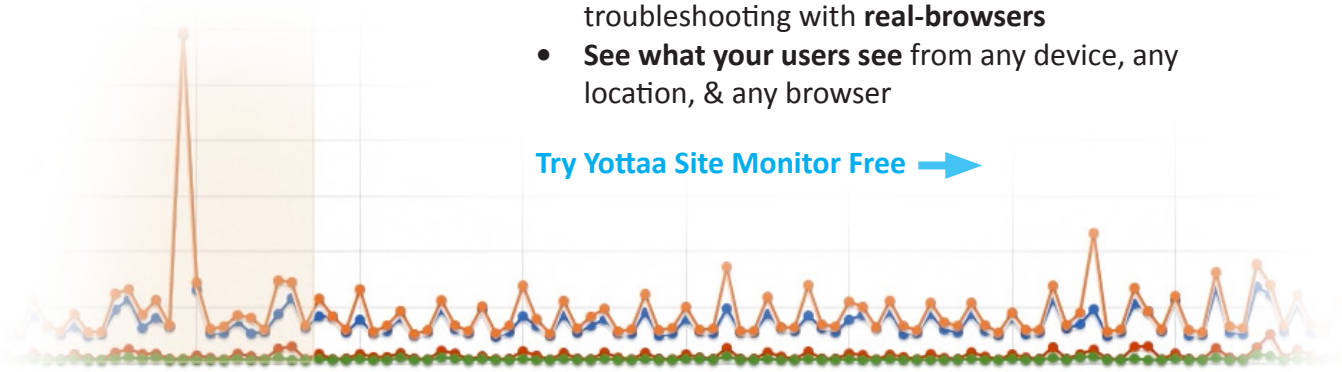

- **24/7 Site Monitoring** is always on alert & ready to help
- Do multivariate testing & web performance troubleshooting with **real-browsers**
- **See what your users see** from any device, any location, & any browser

Try Yottaa Site Monitor Free →

Learn More

Yottaa eBooks

Managing a Web Performance Optimization Project

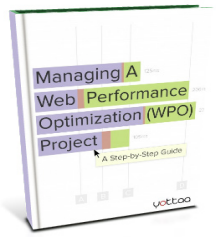

11 Techniques to Make Your Website Rock

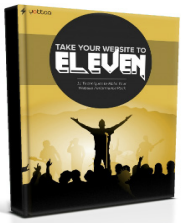

Like this eBook?

Share it!