

10 steps to holiday peak-proof Web Performance

Seasonal peaks raise serious challenges for E-tailers and online businesses to handle traffic surges and stay open.

The average online shopper will wait 4 seconds before going to a competitor. Amazon studies showed a 1% decrease in sales for every 0,1 second increase in response times. The conclusion is that a slow site will hurt sales. Period.

Don't find that out the hard way – Instead apply Web Performance best practices and beat your competition this holiday season.



The holiday peak challenge

Today's web sites are vibrant hubs of content, links and interconnected systems. Content management system and agile development will let you update the site frequently staying in tune with your visitors and their needs. Great! Or is it?

The only problem is that the performance of all those components is seldom tested at their release, and even if it is, they are individually verified, which doesn't mean that the system as a whole is robust. The only complete test you subject your site to is the live test, which happens when your traffic reaches a new peak level. Like the holiday rush.

So instead of putting your combined development efforts to test once a year, in front of your customers, when they need you the most - follow the 10 simple steps to stay ahead of your competition.

1. Align Web Performance goals with business objectives

Too often IT goals are not aligned with, and do not support, business objectives.

Marketing managers might have a goal for new record sales this holiday season, but can your site really handle the load associated with that?

First things first – set business objectives.

Secondly – get business and IT managers together and try to figure out what those business goals mean in terms of peak traffic and loads.

Finally – make sure that your web platform and IT organization are geared up to handle those loads.

2. Be aggressive

Set aggressive performance targets that will not only retain your traffic but also gives you the chance to grab more market share if and when competitor sites are struggling to stay up.

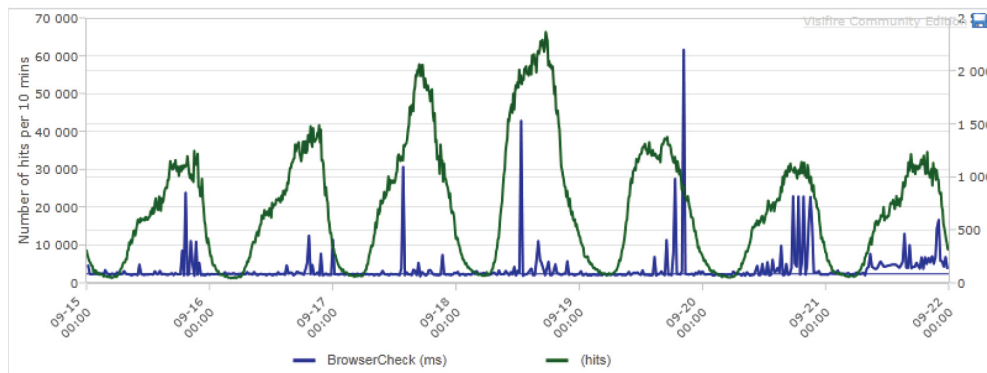
Web shoppers act like stampeding herds in peak times. If one popular site goes down or is slow, fleeing customers will immediately seek out the next alternative increasing the load on that site as well and potentially bringing that one down too.

As a whole web audience is driven towards one or a few sites that manage to stay up, those sites could be grabbing a big chunk of their competitor's customers.

3. Determine maximum performance

At low loads all users on a site have roughly the same response times. At higher loads response times increase in relation to the number of users, and eventually rise exponentially until the site ceases to respond entirely or partly. Establish the maximum performance of the application (i.e. the number of users per minute) with typical user behaviour. This is called a “load test” or a “stress test”, and can be done with a product like Apica LoadTest™.

The graph below shows a real-world example of how response times can vary as the number of simultaneous users go up and down.



4. Take actions towards increased Web Performance

If your load test shows that your Web Performance is not aligned with your business objectives, immediate action must be taken. In order to be able to pinpoint what is causing your performance limit your load testing software must be capable of producing user realistic and analyzable data. Knowing your limit has no real value if you don't know what is causing that limit.

User realistic – A load test should mimic actual user behaviour, such as putting an item in a shopping basket, downloading a file, registering for a newsletter etc.

Analyzable – Data should give a detailed view, not only of the load capacity, but also of the application's behaviour throughout the load test.

5. Action – Verify content size vs. Internet capacity

The total response time is a combination of the individual response times for components on your site. Often a single or a few slow components will cause bad performance. Detailed response time information is necessary to enable an analysis of functions that can be optimized.

A good starting point is to minimize the sizes of images, video clips and the total volume. In particular, verify the sizes of ads, as they often contain a lot of graphics. In summary – Make sure you know what is being shown on your site!

The Internet connection is often a performance bottleneck. Many sites today have very large individual pages: 1-2 MB is not uncommon. That is 10 MB or more during a user session.



If a 1 MB web page is supposed to serve 10 users per second then 80 Mbit per second is needed as the line's "burst volume". If you have a shared connection then the supplier's total capacity may limit your performance. Make sure that you can utilize what you are entitled to via your hosting and/or Internet supplier.

6. Action – Apply and verify front end cache and load balancing

A front end cache improves performance by minimizing backend database calls. If your test analysis shows that the database is a performance bottleneck, this is the way to go. Please refer to the The Apica WebOverload™ product for more information.

Verify that the load balancing is working properly. Load balancers distribute loads from different users onto underlying systems. However, errors due to reconfigurations, for example, can occur. Therefore you should verify that the load balancing works correctly and that the underlying servers receive an evenly distributed load. For example, if one server is at 90% and the other at 10% then you only deliver slightly more than 50% of your total capacity to the end users.

7. Action - Determine performance impact of component interaction

Modern systems consist of many components, such as load dividers, cache servers, web front end and database storage. It is difficult to foresee what will happen at heavy loads; components that function flawlessly at regular loads may all of a sudden become bottlenecks.

Redundancy often creates untested events in itself. Components are often duplicated for better redundancy. But what happens if one or several parts fail? How is performance and stability affected? Can the remaining components carry the load?

8. Action – Control external content and services

The rise of Web 2.0 technologies has made it ever so easy to integrate content, feeds and external functionality into your site. But as you do – their performance bottlenecks now become your headache too.

Not many organisations stop to think of the fact that the combined load time of a page will be determined not only by your own content, but also by that which you fetch from other sites, feeds and APIs.

You should always test and monitor the performance of third party content and have contingency plans if and when their performance starts deteriorating.

9. Action - Manage potential web overload

In addition to determining and increasing your maximum capacity limit, we recommend you add overload protection to your site. An overload protection service does two things:

1. It limits the maximum number of concurrent users on your site so that performance for the allowed users can be guaranteed.
2. It adds extra peak capacity dynamically by a combination of traffic optimization techniques plus cloud-based, on-demand capacity.

This is a sort of “worst case scenario” remedy. If you are uncertain about your peak volumes such a solution can also be a cost effective measure. Fixed cost investments in increased capacity will be idle outside of the peak season, whereas a dynamic capacity solution is only used if and when needed.

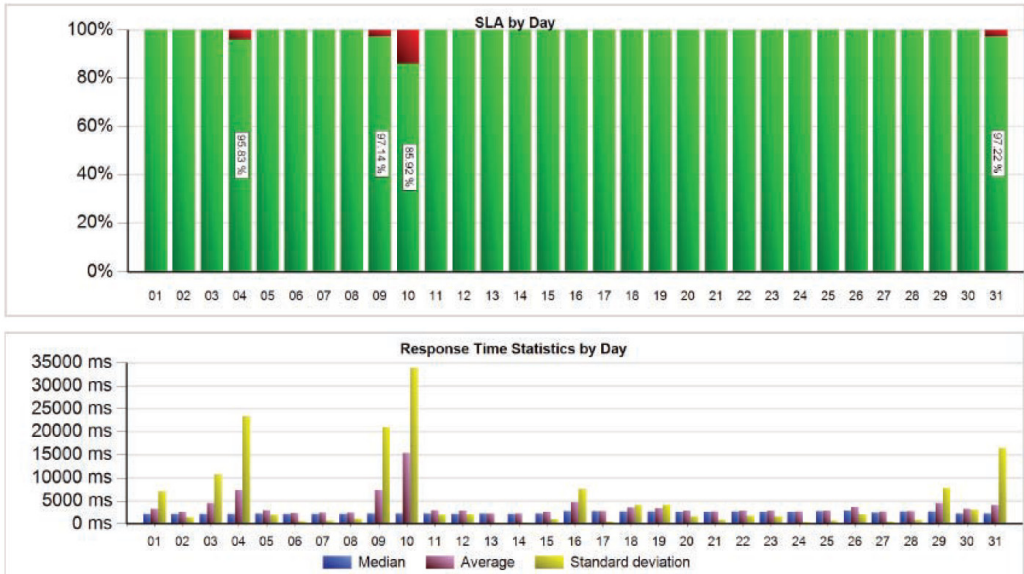
Compare this method to a bouncer that only allows eager shoppers into a physical store as previous customers leave. Otherwise all users will get increasingly poor response times and in the worst case scenario the site ceases to function for all users. It's better to serve the customers who are already on the site and let the others get a polite 'please wait' message. Check out the product sheet for The Apica WebOverload™ controller if you want to know more.

10. Monitor your Web Performance

Ok, so you tested and fixed all the performance bottlenecks, which means you're all set to sit back, relax and enjoy the ride?! Sorry, but no.

During your critical season you should be keeping a close eye on your Web Performance with a product such as the Apica WebPerformance™ monitor. This will let you know the performance of all your critical components in real time and alert you via SMS or e-mail if things go wrong.

You can also set customized SLA targets and monitor your performance vs. those targets. Below you can see an example of how slow response times at peak days impact SLA targets.



Follow the recommendations in this white paper and you will be on your way to a prosperous and trouble free holiday season this year.

Best wishes for a Happy Holiday from the Apica team!

