Peak Season Prep Guide

Preparing your ecommerce site for the next big rush

rackspace®
the open cloud company

Table of Contents

Introduction

Most retailers are familiar with the holiday rush. In brick and mortars, it's the time of year when more hiring happens to ensure smiling faces in the store to help eager customers and enough cashiers at the registers for fast, stress-free checkout. For an ecommerce retailer, beefing up for peak traffic is less about hiring more bodies and more about tuning up the systems and processes that bring, keep, and allow customers to seamlessly make transactions on a site.

Lags in page load not only frustrate users, but **Google's search algorithms** now also penalize sites for slow loads. Take the holiday season — the sales have crept earlier and earlier each year, with some retailers offering deals in October and early November. According to a Shop.org survey of shoppers in 2013, more than 40 percent said they started their shopping in October or earlier. Coupled with the overall growth of online sales, these statistics further increase the need for a well-planned ecommerce strategy, both on the customer-facing side and in your underlying infrastructure.

Getting ready for a peak period, whether it's the holidays, summer tourism, or a far-reaching marketing campaign, involves evaluating both your infrastructure and code to optimize for best results. The cornerstone of peak traffic planning is rigorously load testing your system ahead of time to identify and correct breakpoints and bottlenecks.

1. Test Your Site Before a Traffic Spike

Load testing gives you a window into how your site will perform when peak traffic arrives. The steps below lay out high-level steps involved in load testing.

1. **Determine testing goals:** The initial testing should be a complete end-to-end test. Follow up tests may only cover post-optimization testing of specific pages or processes. Also, keep in mind that different types of users use your site differently. A computer buyer may be divided into segments like consumer, business, and government. The consumer user theoretically takes a different path than a government user, hence a different load on the system that needs optimization. By narrowly defining test objectives and user types, you get more defined results to guide optimization. Use your goals to create a list of questions your test needs to answer, for example:
   - How many concurrent requests can my system handle at maximum load?
   - Are response times for all test paths acceptable?
   - What points in the chain are consuming the most hardware resources?
   - Are there obvious failures caused by large data sets, multiple concurrent users, amount of products on the site, shopping cart functionality, or other factors?
   - Is there any obvious low hanging fruit to optimize? Examples: unnecessary database queries, frequently used code paths that produce a consistent result or frequently repeating database queries that can be cached.
2. **Start with a Benchmark:** Review logs and analytics to see how you've performed during previous peak periods and determine what a typical busy load looks like. Use tools like apache bench or autobench to simulate multiple concurrent users for a benchmark of how many requests per second you're capable of serving. Pay special attention to heavily trafficked pages like home and landing pages where optimization efforts have the biggest payoff. To compare your performance and set benchmarks aligned with other retailers, take a look at **Compuware's Retail Web and Mobile Site Performance Index**.
3. **Determine data collection methods:** Organized data collection is essential to understanding results and learning from them. **LoadRunner** is a software tool that provides sophisticated formatting, flexibility and analysis. **Microsoft Visual Studio** has a SQL script for creating a database repository for results. Evaluate a prospective test data collection tool based on its ability to:
   - Thoroughly document the test conditions
   - Accurately document the results
   - Offer a straightforward analysis of results
   - Archive test data for future comparisons

Use a set of monitoring tools to obtain accurate data about actual test performance in simple visual graphs:

- **Cacti** for capturing metrics
- **MONyog** for monitoring the database
- **statsd** to put stats logging into code to monitor code performance in real time.

4. **Create scripts:** Scripts generate test data and simulate user interaction. The scripts flood the site with requests so that you can identify bottlenecks that occur only during heavy traffic periods. A script will be made for each test path, cookies and all. **JMeter** is an open source load testing software designed to load test functional behavior and measure performance. JMeter scripts are easy to make into templates, and can be copied and pasted to create new ones. Common scripts include:
   - Randomly create products
   - Randomly create orders (fill up a cart, then checkout)
   - Randomly create customer accounts (can combine this with creating an order script)
   - Make changes in the admin interface (creating categories, configuration changes)

These scripts will be run both individually and simultaneously to find performance limits.

5. **Define the Test Environment:** In practice, the load test environment contains:
   - Software tools for sending lots of requests
   - Scripts to simulate user activity on the site
   - Scripts to generate mass amounts of data in the system
   - Hardware to run tools
   - Spreadsheets to track and analyze results

Testing can run for days at a time and use valuable resources. Load testing is a common way to avoid using additional resources to execute tests or degrading other applications during testing. Using the cloud for testing means you can spin up as much CPU as needed to run tests for as long as you need, and then decommission those resources when tests are completed.

6. **Execute Test:** Because running tests can be very time consuming, get the most information possible out of each test run to avoid repeating the same test. To do so, create a spreadsheet with the following information:
   - Name of the person running the test
   - The date, time and duration of the run
   - A clearly defined hypothesis
   - A statement of what has changed
   - A set of metrics being tracked during the test
   - A post-run capture of the results of the metrics, best visualized with a graph of the metric over time

Testing In Action:

One of our customers was gearing up for its first large promotion since establishing its store on our e-commerce software. This promotion included a bulk import of 12,000 coupon codes for specific customers to claim for a one-time purchase.

In preparation, this merchant sent us the 12,000 coupon codes and we ran it through our test environment and quickly found out there would be performance hits to the store cart processes. With the results of this test, we included the development team and optimized the source code that handled coupon codes. By pre-planning with us and testing before sending out these coupon codes we were able to set this campaign up for success

David Blount
IT Director, AmeriCommerce

By saving this information in spreadsheets, you achieve provable results and can identify the most effective tests.

At the completion of the test, you should have a set of data that answers the questions set forth when goals were defined in the beginning. The answers will then guide your next steps for optimization. After optimizing, it's critical to retest to account for any anomalies the optimization may create as fixing one problem can sometimes create another problem that you don't want to wait till you're in the middle of a spike to discover.

2. Solutions for Common Issues that Testing Uncovers

NEED LOAD BALANCING

If site traffic reports reveal an increase in the number of refused connections during load testing, it's time to reassess your load balancing solution. Refused connections are the first sign that your serving capacity is too small for the amount of traffic your site receives. The next sign will probably be visitors calling or emailing to complain that they can't access or transact on your site.

The number of load balancers you deploy is determined by your traffic and performance goals; there is no magic formula. Using the data from the load test will help you determine the number and location of load balancers. Once in place, re-test to confirm that the load balancer can handle the expected load.

NEED BETTER COMPRESSION

Implementing server-side compression can reduce the size of your store pages by reducing the time to return data from the server and making a web server process available sooner. Most modern browsers and web servers are capable of compressing data to send and decompressing at the destination. This helps lower bandwidth requirements but can increase the CPU loads.

Smaller files lead to faster page load times. The two most common options here are to minimize your CSS and JavaScript, removing white space and increasing readability (using a tool like the **YUI Compressor**), and to reduce the size of your images by removing unnecessary data from them with a tool like **smush.it**.

Cloud Load Balancers

On-demand load balancing with SSL termination, session persistence across protocols, advanced node health monitoring and failover protection, and DDoS protection. "With Rackspace Cloud Load Balancing, we get a failover solution without a lot of effort and expense. Our Rackspace Cloud Load Balancer really saves us in terms of the number of machines that we have to manage—and it's all based on technology I know is solid."

H. Wade Minter
Senior Developer + Chief Architect, TeamSnap

NEED A CONTENT DELIVERY NETWORK (CDN)

Using a CDN can help speed up sluggish page loads. On a CDN, the first time content is served to a user, a copy of the content is stored on edge servers geographically closest to that user. Subsequent requests use the stored copy, resulting in faster load time. Hosting landing pages or other heavily trafficked static pages on the CDN helps to maintain a persistent, consistent web presence. As more people visit your site, your landing page will be cached worldwide and as a result load times will improve.

Publishing static content to a CDN rather than the web server can be easily accomplished with services like **Rackspace Cloud Files**, built with Akamai's CDN technology. W3 Total Cache and PressFlow have built-in CDN technology as well.

Cloud Files with CDN

Cloud Files provide you and your applications with online storage for media-items like backups, video files, or user content. You can make this content available to users at high speed over the leading global CDN platform, powered by Akamai Technologies, Inc.

NEED MORE EFFICIENT CODE

There are many services associated with ecommerce and social media that may provide code for you to place either site-wide or on particular pages. These code snippets, used for social sharing, analytics, widgets, etc., are HTML/JavaScript. Many use external requests to the third party, which lag and may not be cached on their end. Where possible, choose solutions for common page elements that avoid third-party code. Another common solution for troublesome JavaScript is to load the code on document ready (when the DOM is fully loaded), rather than including it in your mark-up directly.

Database-heavy sites may benefit from database optimizations, such as schema changes and indexing of commonly queried columns. Avoid using bloated frameworks and libraries.

NEED FULL PAGE CACHING

Caching is an easy way to speed up your application or website which can help with your bounce rate, saving you from potential lost revenue. Determine what data you access frequently and cache it in memory for repeated high-speed access to it. Whether your application is generating static content for web pages or storing sessions in caches, you have to decide how to store those caches. You can store them on your local file system or utilize distributed memory caches like memcached clusters.

NEED AUTOSCALING

Despite your best planning efforts, your site can still experience spikes outside planned peak capacity. For example, a site may typically run on two servers and occasionally traffic spikes up requiring a third server. If that spike is predictable, your team can hop in ahead of time to provision the extra server and decommission it when it's no longer needed.

If that spike is unpredictable, you may be forced to hold on to a third server to cover that spike when or if it happens. With an autoscaling tool, like **Rackspace Cloud Tools Marketplace** partner, **RightScale**, you don't have to hold on to that third server. By setting traffic, performance, or other variables, the third server is launched only when needed and automatically decommissioned when the launch parameters subside. For more auto scaling tools, visit the **Rackspace Cloud Tools Marketplace**.

NEED HYBRID CLOUD

Many still think that if they can't go all cloud, they can't use cloud at all. Not so. A hybrid cloud configuration lets you buy the base and rent the spike. You're able to put the elements in place that you need to run your site now with the ability to burst into the cloud for traffic spikes or to expand functionality without re-architecting or changing platforms.

With **Hybrid Cloud**, an ecommerce store can tap cloud efficiencies for caching, image and video storage, or other resource-intensive, non-critical elements while keeping other elements like payment processing and other security-sensitive site elements on private cloud or on-premises gear to meet PCI compliance and other security.

Caching with Magento:

Magento runs a lot of code to determine which HTML to serve to your client's browser — the full page cache mechanism stores HTML the first time each page is requested and resends that response for all subsequent requests. The page caching feature still varies dynamic content by (e.g., cart count, Welcome message, etc.) while the remaining pieces of the page are served without reprocessing the code.

3. Plan for Peak Success with Rackspace

When lifestyle brand **Alex & Ani** ran an ad for the 2014 Super Bowl, the company knew that site performance was key. "Every second costs money. For every second you add to the loading process, you're reducing your conversion rate by 7 percent, and that adds up," said Ryan Bonifacio, Alex & Ani Vice President of Digital Strategy. The ad brought 28,000 people to Alex & Ani's website and any lag in loading would have been disastrous. Bonifacio came to Rackspace because he was looking for a hosting partner with the necessary performance experience. "Maximizing uptime was one of our biggest requirements in doing diligence with hosting providers. Rackspace was rated the top," he said.

The **Rackspace ecommerce hosting environment** is designed to support customers with resources that effortlessly scale from small to large volumes of traffic. We can help you with the technology needed to develop, test, and scale your site to manage and optimize a spike instead of losing customers.

ENABLING YOUR JOURNEY TO THE NEXT GENERATION OF RETAIL

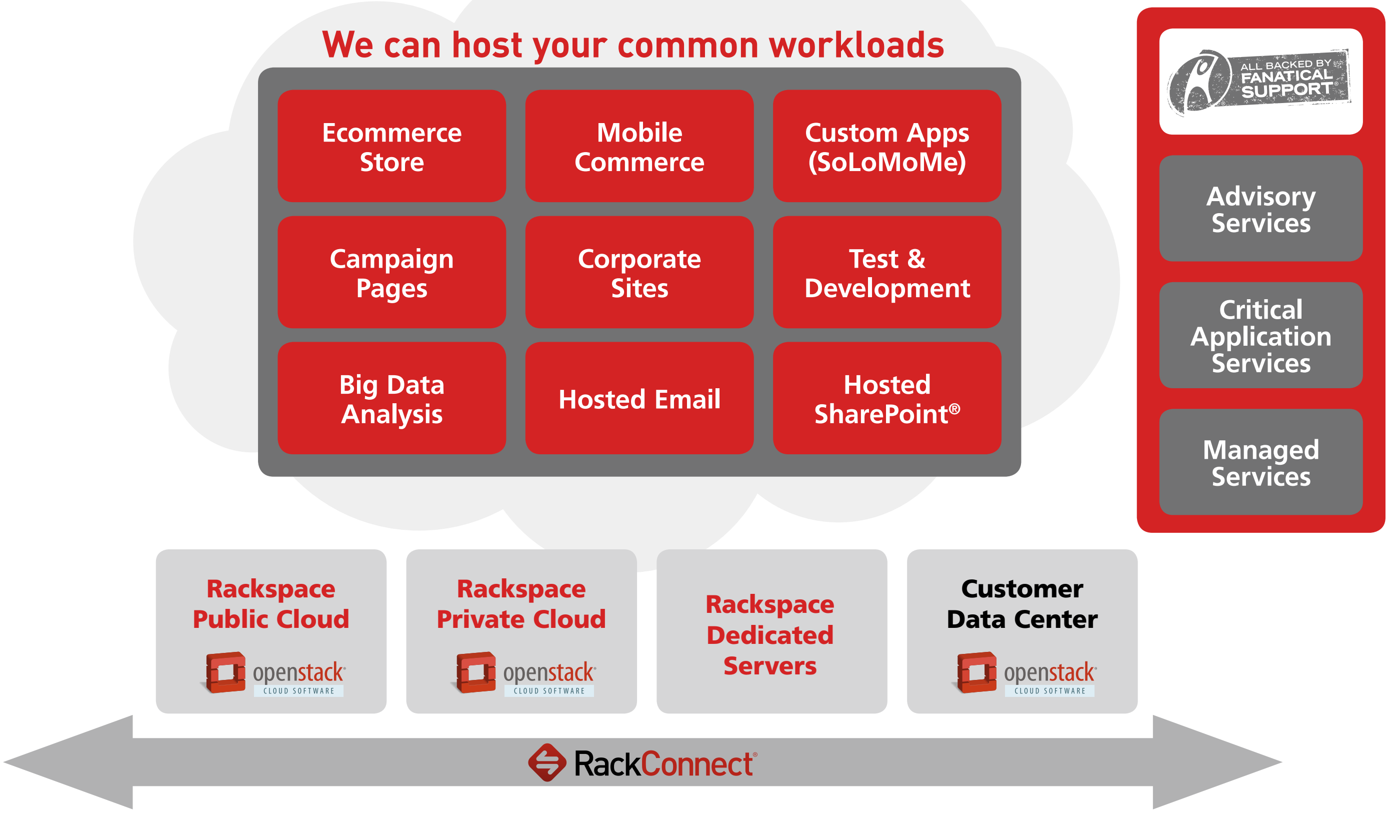

We can host your common workloads

Ecommerce Store

Mobile Commerce

Custom Apps (SoLoMoMe)

Campaign Pages

Corporate Sites

Test & Development

Big Data Analysis

Hosted Email

Hosted SharePoint®

ALL BACKED BY FANATICAL SUPPORT

Advisory Services

Critical Application Services

Managed Services

Rackspace Public Cloud openstack CLOUD SOFTWARE

Rackspace Private Cloud openstack CLOUD SOFTWARE

Rackspace Dedicated Servers

Customer Data Center openstack CLOUD SOFTWARE

RackConnect®

About Rackspace

Rackspace® Hosting (NYSE: RAX) is the open cloud company, delivering open technologies and powering hundreds of thousands of customers worldwide. Rackspace provides its renowned **Fanatical Support**® across a broad portfolio of IT products, including Public Cloud, Private Cloud, Hybrid Hosting and Dedicated Hosting. The company offers choice, flexibility and freedom from vendor lock in.

GLOBAL OFFICES

Headquarters Rackspace, Inc.
5000 Walzem Road | Windcrest, Texas 78218 | 1-800-961-2888 | Intl: +1 210 312 4700
www.**rackspace**.com

UK Office
Rackspace Ltd.
5 Millington Road
Hyde Park Hayes
Middlesex, UB3 4AZ
Phone: 0800-988-0100
Intl: +44 (0)20 8734 2600
www.rackspace.co.uk

Benelux Office
Rackspace Benelux B.V.
Teleportboulevard 110
1043 EJ Amsterdam
Phone: 00800 8899 00 33
Intl: +31 (0)20 753 32 01
www.rackspace.nl

Hong Kong Office
9/F, Cambridge House, Taikoo Place
979 King's Road,
Quarry Bay, Hong Kong
Sales: +852 3752 6465
Support +852 3752 6464
www.rackspace.com.hk

Australia Office
Level 4, 210 George Street,
Sydney, NSW 2000
Phone: 1-800-722577
www.rackspace.com.au

This whitepaper is for informational purposes only. The information contained in this document represents the current view on the issues discussed as of the date of publication and is provided "AS IS." RACKSPACE MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, AS TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS DOCUMENT AND RESERVES THE RIGHT TO MAKE CHANGES TO SPECIFICATIONS AND PRODUCT/SERVICES DESCRIPTION AT ANY TIME WITHOUT NOTICE. USERS MUST TAKE FULL RESPONSIBILITY FOR APPLICATION OF ANY SERVICES AND/OR PROCESSES MENTIONED HEREIN. EXCEPT AS SET FORTH IN RACKSPACE GENERAL TERMS AND CONDITIONS, CLOUD TERMS OF SERVICE AND/OR OTHER AGREEMENT YOU SIGN WITH RACKSPACE, RACKSPACE ASSUMES NO LIABILITY WHATSOEVER, AND DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO ITS SERVICES INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT.

Except as expressly provided in any written license agreement from Rackspace, the furnishing of this document does not give you any license to patents, trademarks, copyrights, or other intellectual property.

Rackspace, Fanatical Support, and/or other Rackspace marks mentioned in this document are either registered service marks or service marks of Rackspace US, Inc. in the United States and/or other countries. OpenStack is either a registered trademark or trademark of OpenStack, LLC in the United States and/or other countries. Third-party trademarks and tradenames appearing in this document are the property of their respective owners. Such third-party trademarks have been printed in caps or initial caps and are used for referential purposes only. We do not intend our use or display of other companies' tradenames, trademarks, or service marks to imply a relationship with, or endorsement or sponsorship of us by, these other companies.

References:

1 https://www.internetretailer.com/2013/02/04/unwrapping-holidays