



CITO Research

Advancing the craft of technology leadership

The Secret to Business Results: How Network Connections Drive Web Application Performance



Sponsored by

EQUINIX



Contents

Executive Summary	1
Introduction	1
Optimizing the Application Architecture	2
Renting Scalability: The Public Cloud	3
Getting Faster and Reducing Cost	3
Close to the In-Crowd: Speeding Up Applications	4
Finding the Right Caching and Distribution Strategies	5
Distributing Data Centers: Getting Up Close and Personal with Users	6
Using Commercial Cache Arrays or CDNs: Pushing Content to a Public Edge	7
Using Private CDNs: Pushing Content to a Custom Edge	7
Challenges to Network Quality	8
Challenge: Getting Burned by Hot Potato Routing	8
Challenge: Delivering Availability	9
Challenge: Delivering Consistency	9
Conclusion	10



Optimization of web applications is most effectively achieved by taking advantage of the advanced connectivity available at carrier-neutral data centers.

The Secret to Business Results: How Network Connections Drive Web Application Performance

Executive Summary

The application delivery model has changed. The default paradigm of an application has evolved from users sitting at desktops attached to the corporate data center to mobile access via phones and tablets, each requiring its own optimized experience. Add to that the fact that both users and application components are literally all over the world and that users expect web applications to act like desktop applications—with immediate response and features like copy and paste, drag and drop, highlighting, and more. The challenge is delivering applications effectively, regardless of where the users are and what device they have in their hands.

Different applications also have wildly varying performance requirements. For some Software-as-a-Service applications, latency of 150 milliseconds is perfectly acceptable while real-time interactive applications like multi-player games can tolerate far less. Application design plays a critical role, but optimizing networks is at least as important in reducing latency.

In this white paper, CITO Research discusses the challenges of delivering applications effectively when your users and applications are global in scope. It explains how you can meet those challenges and the expectations of your users.

Introduction

Some of you remember the good old days—when it was easy to make your web application fast. Everything was inside your data center. One pipe brought traffic in from the public Internet or from other corporate locations, and the front-end servers, the app servers, and the database were all optimized to work together to deliver as much speed and scale as possible.

Those days are gone.

Today, applications are not on a LAN, and users aren't coming in from one pipe. Both the application components and the users are spread all over the world. Today, making your web application fast means that you need fast network connections between all the components of that application.

CITO Research has found that optimization of web applications is most effectively achieved by taking advantage of the advanced connectivity available at carrier-neutral data centers run by companies such as Equinix.



Using carrier-neutral data centers, Badgeville reduced infrastructure costs by 40%.

The Secret to Business Results: How Network Connections Drive Web Application Performance

If a website is slow, consumers are out of there. According to Equation Research, consumers won't wait:

- 32% *abandon* sites after one to five seconds
- 39% cite *speed* as more important than *functionality*
- 37% say they *won't go back* to a slow site
- If a site is slow, 27% will *switch to a competitor's site*

To meet consumers' need for speed, many of the top Internet companies connect to as many ISPs as they can through locations known as carrier-neutral data centers. The results are dramatic. [Badgeville](#), for example, was able to redesign applications to reduce infrastructure costs by 40%. [Box.com](#) was able to dramatically increase levels of service to its rapidly growing international client base and improve customer application response times by 60%.

The bottom line is that network optimization of web applications offers a better, more consistent user experience and faster, more reliable network performance.

Network Optimization of Web Applications

Network optimization of web applications involves situating critical infrastructure close to telecom providers, cloud service providers, content distribution networks, and application component providers. The object is to place assets as close as possible to the users, with the lowest latency possible, to maximize application performance.

Optimizing the Application Architecture

Today's web applications are composed of a combination of services from a wide collection of Internet service partners. A single page on a major website may be comprised of as many as 100 scripts, trackers, embedded pages, and dynamic objects.¹ The key to optimizing web applications for both performance and cost is to make sure you and your customers can get to the services needed and make each of them work at the best possible price/performance tradeoff.

¹ According to [WebSiteOptimization.com](#), as of November 2012, the 1000 top web pages have about 100 objects each, nearly doubling from an average of 49.9 objects in January 2009.



Network optimization opens the door to both performance and cost savings.

The Secret to Business Results: How Network Connections Drive Web Application Performance

Renting Scalability: The Public Cloud

Ecommerce applications are subject to unpredictable demand spikes, and numerous factors can exacerbate this problem, such as a sudden unexpected increase in activity, the kind of thing that happens if a celebrity casually mentions your website. When this happens, ecommerce applications need to be able to rapidly and dynamically scale to meet demand. Many top businesses are designed to grow from their baseline deployment into the vast resources of the public cloud during times of high demand, sometimes called “bursting into the cloud.” To retain customers, your apps have to scale and consistently deliver excellent performance.

To do this effectively, some cloud users place non-cloud infrastructure in close proximity to public cloud providers like Amazon Web Services (AWS). At Equinix, a customer can get a direct fiber or copper connection to AWS via AWS Direct Connect. An Equinix carrier-neutral data center brings together many cloud ecosystem players offering similar direct connection options.

Sometimes it's not compute power but storage that is needed. Distributed storage providers such as Nirvanix provide direct connections to their service at carrier-neutral data centers. With a direct connection, organizations can connect their corporate backbone to their storage provider and increase the security and performance of these services for their end users. Additionally, when using a direct cross-connect, Storage-as-a-Service solutions can be kept segmented from the public Internet to achieve security, performance, control, and cost benefits.

If an ecommerce company hosts all of its websites in the public cloud, but needs to manage its user credentials and credit card information in a more secure setting, the company no longer has to sacrifice retrieval speed by housing that data in a data center miles away from the cloud. A direct connection between managed storage assets and public cloud computing power in a carrier-neutral data center meets security requirements while providing less than 10 milliseconds of latency.

Getting Faster and Reducing Cost

Network optimization opens the door to both performance and cost savings.

For example, the Badgeville Behavior Platform is a public-cloud-based solution to help businesses influence and motivate end user behavior by providing goals and rewards for completing tasks. As the company began to grow, it ran into challenges keeping up with demand for its services because it could not quickly and easily scale to meet the bandwidth and computational requirements of its success. Badgeville wanted to easily leverage the public cloud to achieve



The Secret to Business Results: How Network Connections Drive Web Application Performance

economies of scale on commodity hardware, but found that its current solution made that difficult. In addition, the public cloud didn't offer the specific database capabilities they needed. Badgeville needed a combination (or hybrid) cloud solution.

To meet their performance goals and achieve ease of scaling in the future, Badgeville first migrated its databases from the public cloud to a private deployment in an Equinix location. Next it established a gigabit low-latency direct connection between its database servers and its public cloud provider inside Equinix. The direct connection allowed Badgeville to connect its private deployment and public cloud applications to get all the benefits of the public cloud while maintaining its database outside of the public cloud—an optimal combination of technologies into a high performance solution.

The best part is that the new approach saved Badgeville over 40% in infrastructure costs. Before moving the database servers to Equinix, Badgeville ran its database using Amazon High-Memory Quadruple Extra Large nodes via the Engine Yard Platform-as-a-Service cloud. Running each of these database nodes cost about \$2,000 per month. The two new servers that Badgeville put into the Equinix facility cost only \$4,000, so they paid for themselves in just two months.

Because of the performance gains from going to private database servers, Badgeville now has significantly more capacity for growth with its existing footprint.

Close to the In-Crowd: Speeding Up Applications

Reports from some of the top ecommerce companies show that a slow and unresponsive user experience is one of the leading factors in abandoned shopping carts and lost sales. Today's ecommerce companies must be able to quickly identify the end user, pull targeted content, and deliver an optimized user experience. This sequence of events relies on quick and reliable performance from many companies working in close coordination with one another.

Many of the top-tier Internet companies have already defined a network strategy around carrier-neutral data centers because they get the fastest access to business partners and critical network infrastructure there. In many of the low-latency financial trading markets, the companies that connect to the financial exchange all reside within the same data centers and use a direct connection to ensure the lowest possible latency and the best chance to place the winning bid before anyone else. Real-time bidding for advertising is also very latency sensitive and companies are migrating to the same type of direct-connect solution. In fact, the benefits of being able to directly access many network and service providers are

A slow and unresponsive user experience is one of the leading factors in abandoned shopping carts and lost sales.



There's only one way to reduce network latency: putting your application at an advantageous location.

The Secret to Business Results: How Network Connections Drive Web Application Performance

so great that many other industries choose to deploy in the same well-connected data centers. The result is that when these assets are located in close proximity, application performance is enhanced, unexpected issues decrease, and both revenue and end user satisfaction go up. Choosing a location that has the services your applications need offers a huge boost in reliability and performance.

Finding the Right Caching and Distribution Strategies

When application performance deteriorates under heavy transaction volume or when performance is degraded by network latency, engineers often turn to vendors that offer solutions to compress transmitted data. With some solutions, such as WAN acceleration, two or more endpoints compress the data at the source and decompress it at the destination. With others, such as web optimization, the content is optimized in-stream (for example, image, HTML, and CSS size reduction) and doesn't require any receiving end reassembly. Both techniques help organizations operate more effectively over carrier connectivity.

However, compression and reduction techniques only impact the content size and don't affect the latency between the end user and the web application. There's only one way to reduce network latency: putting your application at an advantageous location. When serving geographically distributed users, organizations can improve performance and reliability by serving content closer to the end user. The original models of web serving placed a tier of web servers in a central location, which any number of users would access. The greater the network distance to end users and the greater the volume of data, the poorer the application performance. Typically, organizations in this position use caching—placing and updating copies of the most accessed components close to the end user.

Caching is usually approached in one of three ways, as shown in the following table and described in more detail in the following sections.



*The Secret to Business Results:
How Network Connections Drive Web Application Performance*

Three Approaches to Caching

Cache Approach	Transmission of Data	Number of Data Centers	Data Center Location	Cost Model	Long-Term Cost
Build data centers in major regions	Leased telecom lines from nearest Central Office to Data Center	Multiple large deployments	Geographically deployed	Data center build cost, support cost, (# of telcos) x (mile distance) x (bandwidth size)	High
Public CDN	Commercial cache arrays	One	Centralized	Pay on a per-meg basis to deliver content to each user	Medium
Private CDN deployment in multiple carrier-neutral facilities	Peering, zero-mile transit, paid peering for direct access to MSOs	Many small deployments in many major metro areas	Widely distributed according to use	Directly connect to peers, purchase transit from a variety of providers	Low

**Distributing Data Centers:
Getting Up Close and Personal with Users**

One approach to getting better performance is to decentralize the web application and deploy distributed assets in multiple data centers chosen strategically to be close to users. Bechtel, the international construction contractor, used a leased line to connect its main data centers with smaller data centers in far-flung localities. Bechtel decided to take a page from web scale companies that distribute their infrastructure at multiple well connected points on the globe and leverage network dense connectivity points for performance and cost benefit.

Bechtel virtualized a large portion of its equipment and placed it in Equinix's data centers (Equinix has carrier-neutral data centers around the world).

This allowed Bechtel to switch their strategy from one large carrier of leased lines between each data center to directly connecting to their choice of carriers in each market. This move saved almost 40% of its communications costs, reduced latency, and placed it in a location where Bechtel has low-latency access to an ecosystem of cloud service providers such as Workday and [Salesforce.com](https://www.salesforce.com).



Box's product is 2.7 times faster than its top competitor.

The Secret to Business Results: How Network Connections Drive Web Application Performance

Box.com used the same strategy to enhance its service offering. As part of a feature called Box Accelerator, Box was able to leverage Equinix's global data center presence to offer direct connect access to their largest customers for improved performance. As a result, according to third-party testing provider Neustar, Box Accelerator provides the lowest average upload times across all locations tested. Box's product is 2.7 times faster than its top competitor.²

Using Commercial Cache Arrays or CDNs: Pushing Content to a Public Edge

Many organizations use commercial caching services to distribute applications to regional markets. In this case, a portion of the customer's web application assets is mirrored on a shared network of cache servers which then serve web pages to end users, saving the customer the expense of building out its own deployments to handle the web application demand.

Content delivery networks (CDNs) provide these cache services to most commercial businesses. In some cases, the customer pre-loads the assets to the cache provider, which then takes over distribution to local markets. If an asset is frequently updated, the new content has to be continuously reloaded so some firms position their application servers at carrier-neutral data centers operated by firms such as Equinix to get higher speed, increased reliability, and lower cost access to CDN providers.

Using Private CDNs: Pushing Content to a Custom Edge

The largest bandwidth consuming web applications, such as those that stream high-definition video, have such high costs of distribution and optimization that many organizations are looking for a way to take that distribution in-house. To do this effectively, those organizations must have enough bandwidth and servers deployed to serve all the content to their users. Leveraging carrier-neutral data centers, such as those offered by Equinix, facilitates access to a wide range of network providers to enable quick scalability of bandwidth and ease the distribution of content to users.

² <http://blog.box.com/2012/09/moving-at-the-speed-of-business-introducing-box-accelerator/>



Delivering consistent, reliable service is fundamental to customer conversion and retention.

The Secret to Business Results: How Network Connections Drive Web Application Performance

Challenges to Network Quality

To reduce overall cost, network providers come together at peering points to efficiently exchange traffic with a multitude of other providers all in one place. By deploying in these same locations, organizations can enjoy the benefits of being able to directly and easily access these same providers. Equinix hosts many of the primary peering locations in each region for these carriers. By colocating at the right data centers and being able to directly access multiple ISPs' backbones, companies can achieve significant performance benefits that allow the major companies to be fast, independent of their users' locations. This not only affects user experience but also revenue.

Performance isn't just about the speed of a site; availability and consistency are also important. Delivering consistent, reliable service is fundamental to customer conversion and retention.

Challenge: Getting Burned by Hot Potato Routing

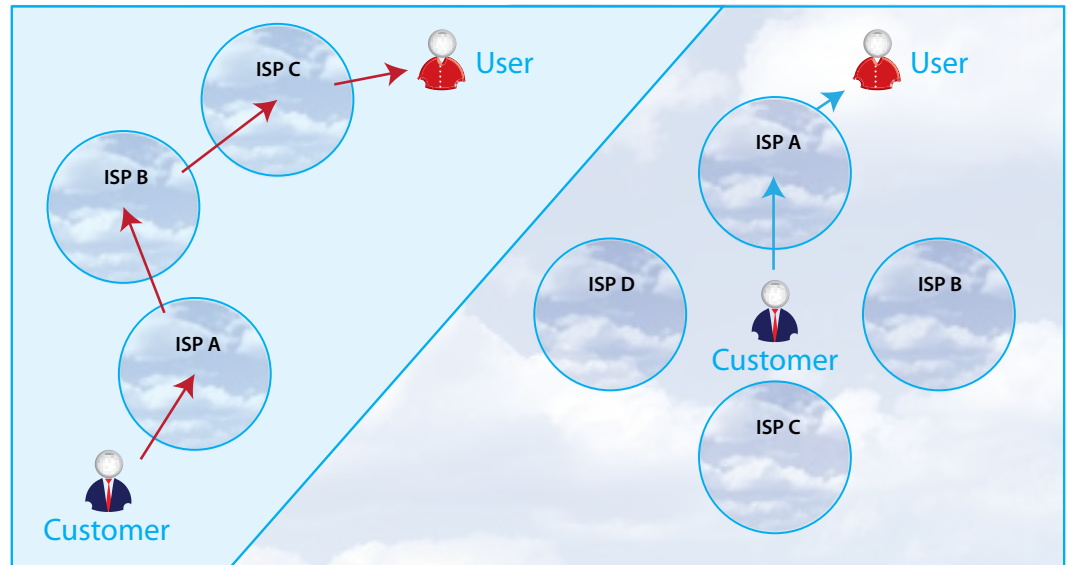
Some people are surprised to learn that network carriers are not responsible for providing the highest performing link between an application and its customers. Almost all carriers perform "hot potato routing," which means the carrier will hand off a connection to another provider at the first, most directly available (and least expensive) opportunity. This can result in suboptimal routing and performance degradation. Organizations can take advantage of carrier-neutral data centers to avoid the impact of hot potato routing.

For example, if a customer is in Japan, it is best to connect directly to a Japanese Internet service provider (ISP) instead of one that routes your traffic the long way from the US to Europe to Japan. By connecting application servers directly to the backbone of regional carriers, such as KVH in the Netherlands or NTT in Japan, application performance can be improved by up to 30%.

Another example is customers who may access your application from home or from a remote office. These locations may use a local ISP or multi-service operator (MSO) that lacks a local peering connection to your carrier. Data between the user and the application has to travel via your carrier to the location of the application server. The closer the application server is to local customers—ideally where their ISPs or MSOs connect to the global provider's backbone—the better the application performance.



*The Secret to Business Results:
How Network Connections Drive Web Application Performance*



Hot potato routing versus use of carrier-neutral data centers

Challenge: Delivering Availability

Using a data center provider that also serves as a carrier-neutral interconnection point can have a huge impact on availability, that is, the ability to reach an application from anywhere on the Internet. By using Equinix data centers to widely connect to many different carriers, companies find that their lack of availability due to network issues can drop from hours a year to zero.³

Challenge: Delivering Consistency

Regardless of where customers are in the world, and no matter what kind of device they're using, the ubiquity of the Internet has led them to demand a consistent level of performance and availability. In delivering applications, providers must manage many interconnections and peering points. The greater the number of network hops in a path, the greater the potential for inconsistencies that can mar a smooth application experience.

³ <http://blog.equinix.com/wp-content/uploads/2011/04/Equinix-Optimizing-Internet-Application-Performance.pdf>



To win in the web application market, providers are challenged to increase application performance and availability while simultaneously reducing costs.

This paper was created by CITO Research and sponsored by Equinix.

The Secret to Business Results: How Network Connections Drive Web Application Performance

Conclusion

To win in the web application market, providers are challenged to increase application performance and availability while simultaneously reducing costs. Global application deployment, acceleration, and scaling are fundamental requirements for any organization that wants to compete in the digital marketplace.

Carrier-neutral data centers provide access to many of the resources needed to help application performance meet and exceed expectations, reduce cost and complexity, and capture revenue. Delivering content between these points using content delivery networks increases the efficiency of data transfer. The next white paper in this series, "Perform or Die: Navigating the Content Distribution Marketplace," describes forward-thinking strategies for efficient and cost-effective content delivery using private CDNs.

About Equinix

Equinix, Inc.

One Lagoon Drive

4th Floor

Redwood City, CA 94065

Main: +1.650.598.6000

Fax: +1.650.598.6900

Email: info@equinix.com

CITO Research

CITO Research is a source of news, analysis, research, and knowledge for CIOs, CTOs, and other IT and business professionals. CITO Research engages in a dialogue with its audience to capture technology trends that are harvested, analyzed, and communicated in a sophisticated way to help practitioners solve difficult business problems.

Visit us at <http://www.citoresearch.com>